Linköping studies in science and technology. Dissertations.
No. 1520

# Autonomous Localization in Unknown Environments

**Jonas Callmer**


Department of Electrical Engineering
Linköping University, SE–581 83 Linköping, Sweden

Linköping 2013

**Cover illustration:** Example of positioning result using a foot-mounted inertial magnetic sensor unit. The user is visiting every room in a corridor, similar to how fire fighters operate. The result is overlaid on a floor plan that was unknown to the positioning system.

Linköping studies in science and technology. Dissertations.
No. 1520

**Autonomous Localization in Unknown Environments**

Jonas Callmer

*callmer@isy.liu.se*
*www.control.isy.liu.se*
*Division of Automatic Control*
*Department of Electrical Engineering*
*Linköping University*
*SE–581 83 Linköping*
*Sweden*

ISBN 978-91-7519-620-6 ISSN 0345-7524



Printed by LiU-Tryck, Linköping, Sweden 2013

*Till mina föräldrar*

## Abstract

Over the last 20 years, navigation has almost become synonymous with satellite positioning, e.g. the Global Positioning System (GPS). On land, sea or in the air, on the road or in a city, knowing ones position is a question of getting a clear line of sight to four or more satellites. Unfortunately, since the signals are extremely weak there are environments the GPS signals cannot reach but where positioning is still highly desired, such as indoors and underwater. Also, because the signals are so weak, GPS is vulnerable to jamming. This thesis is about alternative means of positioning for three scenarios where GPS cannot be used.

Indoors, there is a desire to accurately position first responders, police officers and soldiers. This could make their work both safer and more efficient. In this thesis, an inertial navigation system using a foot mounted inertial magnetic measurement unit is studied. For such systems, zero velocity updates can be used to significantly reduce the drift in distance travelled. Unfortunately, the estimated direction of movement is also subject to drift, causing large positioning errors. We have therefore chosen to throughly study the key problem of robustly estimating heading indoors.

To measure heading, magnetic field measurements can be used as a compass. Unfortunately, they are often disturbed indoors making them unreliable. For estimation support, the turn rate of the sensor can be measured by a gyro but such sensors often have bias problems. In this work, we present two different approaches to estimate heading despite these shortcomings. Our first system uses a Kalman filter bank that recursively estimates if the magnetic readings are disturbed or undisturbed. Our second approach estimates the entire history of headings at once, by matching the cumulative sum of gyro measurements to a vector of magnetic heading measurements. Large scale experiments are used to evaluate both methods. When the heading estimation is incorporated into our positioning system, experiments show that positioning errors are reduced significantly. We also present a probabilistic stand still detection framework based on accelerometer and gyro measurements.

The second and third problems studied are both maritime. Naval navigation systems are today heavily dependent on GPS. Since GPS is easily jammed, the vessels are vulnerable in critical situations. In this work we describe a RADAR based backup positioning system to be used in case of GPS failure. RADAR scans are matched using visual features to detect how the surroundings have changed, thereby describing how the vessel has moved. Finally, we study the problem of underwater positioning, an environment GPS signals cannot reach. A sensor network can track vessels using acoustics and the disturbances they induce to the earth magnetic field. But in order to do so, the sensors themselves must first be positioned. We present a system that accurately positions underwater sensors using a friendly vessel with a known magnetic signature and trajectory. Simulations show that by studying the magnetic disturbances caused by the vessel, the location of each sensor can be accurately estimated.

## Populärvetenskaplig sammanfattning

För att bestämma sin position utomhus använder man idag ofta satellitbaserade positioneringssystem såsom det amerikanska GPS-systemet. På en öppen plats kan man med en enkel mottagare få reda på sin position på några meter när inom loppet av ett par sekunder. Det krävs att platsen är öppen eftersom man måste kunna se minst fyra satelliter samtidigt för att positionen ska kunna bestämmas. I de flesta fall är det inget problem, men i vissa fall innebär kravet en stor begränsning.

Den här avhandlingen handlar om problemet att kunna bestämma sin position i miljöer där GPS inte går att använda. Vi har valt att studera positionering i huvudsak inomhus men även på och under vatten. Signalerna från GPS-systemet är nämligen så svaga att de inte kan penetrera byggnader eller vatten alls. I de få fall man kan ta emot en signal, har den oftast studsat mot andra objekt på vägen. Eftersom man då inte har fri sikt till fyra satelliter så går det ändå inte att positionera sig med god precision. För att positionera sig i sådana miljöer krävs det alltså helt nya metoder.

Idag efterfrågas inomhuspositioneringssystem för exempelvis brandmän, soldater och poliser. Tanken är exempelvis att om ett rökdykarpars positioner är kända i realtid kan räddningsinsatsen bli både säkrare och effektivare. Deras uniformer kan innehålla specialanpassad utrustning som räknar ut rökdykarnas positioner och vidareförmedlar dessa till berörda parter.

Rökdykare arbetar ofta i okända inomhusmiljöer. I framtiden kan man tänka sig att digitala kartor kommer att finnas tillgängligt för räddningstjänsten om det börjar brinna i till exempel ett hotell, ett sjukhus eller i en skola. Kartan kan då användas för att hjälpa till att beräkna rökdykarens position i en stor byggnad. I den här avhandlingen har vi studerat problemet att positionera sig i miljöer där kartor troligtvis inte kommer finnas tillgängliga under många år framöver, såsom i lägenheter, villor, små kontor och dylikt.

Positioneringssystemet baseras på en sensor som mäter acceleration, vinkelhastighet och magnetfält. Vinkelhastighet är hur snabbt sensorn roterar. Acceleration tillsammans med vinkelhastighet kan användas för att uppskatta hur användaren rör sig och magnetfältet kan ge information om hur användaren är riktad i förhållande till norr. Systemet är ett så kallat dödräkningssystem eftersom det utgår från en position och sedan försöker beräkna förflyttningen från denna punkt. Sensorn är fastsatt på foten, och användarens rörelser beräknas steg för steg. Genom att utnyttja att sensorn är still när foten är i kontakt med marken så kan ett stegs längd beräknas på några centimeter när. Dödräkningssystemet baseras alltså på steg för steg-förflyttningar som sker i antagna riktningar, men om rökdykaren egentligen rör sig i en annan riktning än den man tror, kommer den uppskattade positionen bli väldigt fel väldigt snabbt. Att ha en mycket god kännedom om i vilken riktning rökdykaren faktiskt rör sig är således avgörande.

Det problem som studeras mest i den här avhandlingen är därför riktningsskatt-

ning. Vi använder en magnetfältssensor som mäter jordens magnetfält och därför ger oss riktningen mot norr, samt ett gyro som mäter vinkelhastigheter för att beräkna riktningen. Tyvärr är det dock svårt att mäta jordens magnetfält inomhus. Stålstrukturer, elskåp, kablar och liknande producerar magnetiska störningar som ofta är starkare än jordens magnetfält. Sådana störningar är mycket vanliga inomhus varför magnetfältssensorn är mycket opålitlig inomhus. I den här avhandlingen presenteras två metoder att uppskatta användarens riktning trots dessa störningar. Dessa metoder har inkluderats i dödräkningssystemet för att minska de fel som orsakades av osäker riktning. Experiment visar att felet i den uppskattade positionen minskar betydligt när pålitlig information om användarens riktning finns tillgängligt.

Eftersom signalerna från GPS-satelliterna är så otroligt svaga så är de också mycket känsliga för störningar. Idag är det lätt att införskaffa störsändare som sänder ut signaler som överröstar signalerna från satelliterna, vilket slår ut all GPS-positionering inom ett stort område. Ofta är det dock inte illvilja som ligger bakom sådana störningar. I de flesta fall är det användaren själv som omedvetet sänder ut signaler inom fel frekvensband som förstör positioneringen. I somliga miljöer, såsom fartygsnavigation, används GPS av många olika system varför en störning i fel läge kan få allvarliga konsekvenser. Vi har därför tagit fram ett reservsystem för positionering som är oberoende av GPS. Systemet är istället baserat på fartygets egen RADAR. Ett sådant system kan användas för att detektera GPS-störningar och minska dess inverkningar.

Positioneringssystemet använder sig av de bilder av omgivningen som RADAR-systemet ger. Öar och kustlinjen avtecknas väl och avståndet och vinkeln till dessa kan mätas med god precision. Ofta finns det vissa objekt i omgivningarna som syns särskilt väl såsom hus eller branta klippor. Sådana objekt kan observeras över en lång tid och kallas för landmärken. Genom att studera flera sådana landmärken och hur deras positioner relativt fartyget förändras över tiden, kan fartygets förflyttning beräknas. Vi presenterar ett positioneringssystem som använder sådana RADAR-baserade landmärken och experiment visar att fartyget kan positioneras med god precision över långa sträckor.

Det tredje problemet som behandlas i avhandlingen är undervattenspositionering. Att exakt kunna bestämma en position under vatten är generellt sett dyrt och svårt eftersom GPS-signalerna inte når ner. I det här fallet är det speciella undervattenssensorer som vi ska bestämma positionen för, varför problemet går att lösa med hjälp av sensorernas egna mätningar.

För att upptäcka och följa inkommande fartyg kan man använda ett nätverk av sensorer på havets botten. Sensorerna känner av de magnetfältstörningar som fartyget ger upphov till och uppfattar även ljudet från fartygets framdrivningssystem. Dessa kan användas för att positionsbestämma ett fartyg med god precision, givet att sensorernas egna positioner är kända. I den här avhandlingen studerar vi problemet att positionsbestämma sensorerna själva med hjälp av ett fartyg som ger upphov till en känd magnetisk störning. Fartyget framförs genom området där sensorerna finns och de magnetiska störningar som fartyget ger upphov till

kan användas för att beräkna sensorernas positioner. Om även GPS-mätningar av fartygets rutt tas med i beräkningarna så kan sensorernas positioner bestämmas med god precision. Systemet utvärderas med hjälp av simuleringar.

## Acknowledgments

This thesis marks the end of a journey. Five years went by pretty fast even though it doesn't feel that way when one is in the middle of it. The improvements one experience are slow but steady until one day you realize that you have actually learned quite a lot.

The person who has been my greatest source of learning is my supervisor Professor Fredrik Gustafsson. Your broad knowledge and experience has been a great guiding light throughout this journey. Thank you for your never ending ability to find time to assist despite your extremely busy schedule. The final proof to me of my progress was when we went from you always telling me 'Isn't that what I said a month ago?' to me finally getting to tell you once: 'That's what I said a month ago'. It has been a pleasure.

For all those things I needed to discuss, my co-supervisor Dr David Törnqvist has always been there to help. Your enthusiastic and helpful manners combined with your solid knowledge makes me coming back to learn more. Thank you for your assistance!

The Automatic Control group is a very well functioning group with a welcoming and helpful atmosphere combined with a culture of hard work and a willingness to excel in all fields. It is steadily steered by Professor Svante Gunnarsson who took over after Professor Lennart Ljung who built the group from the ground up. Not until I got a good insight into how some other groups function, did I fully realize how exceptionally well organized our group really is. You have both done a tremendous job.

Whenever I run into administrative issues, our secretary Ninna Stensgård and before her Åsa Karmelind and before her Ulla Salaneck, have always been there to help. Thank you!

I would also like to thank Dr Gustaf Hendeby and Dr Henrik Tidefelt for creating an astounding LaTeXtemplate making writing this thesis a matter of only that: writing. Your attention to detail never ceases to amaze me. Various parts of this thesis have been proofread by Fredrik, David and Dr Karl Granström. Your input has significantly improved this thesis and I thank you for that. All remaining errors are naturally mine.

A little more than a year ago, I had the good fortune of doing a predoc in the US. When the original plans were cut extremely short, Professor Arogyaswami Paulraj at Stanford University was kind enough to welcome me to visit his group instead. While there, Professor Stephen Boyd took time from his intense schedule to supervise me. Your suggestion that I should look at the heading estimation problem spawned plenty of ideas, that now constitute about half of this thesis. Thank you very much for your time and enthusiasm! I would also like to thank Professor Peter Stoica who helped to set up this arrangement on such a short notice. The research school Forum Securitatis of which I have been participating,

deserve a special acknowledgement for funding the visit. It was a great experience for me and I want to thank you all so very much.

During these years I have been lucky enough to have some of my best friends as colleagues.

Dr Karl Granström and I go back to the days of our master thesis project and despite seeing you almost every day for something like seven years, I do not get tired of your company. We have had many good times in Sydney, Tasmania, Samoa, Japan, the Balkans and of course in Linköping and even though this marks the end of an era I do not think it marks the end of a great friendship.

Lic Martin Skoglund and I go back even longer, sharing a few thousand cups of coffee and thereby plenty of time to look for treasure. Enough time it turned out, to make Maria jealous of me more than once. Thank you for the happy times that have been and the ones that are yet to come.

Dr Christian Lundquist is a man of many talents and his great sense of humour always makes him terrific company. Our discussions have been plenty and I have enjoyed them all. And I will never forget your comment 'I don't like nature!' halfway through our trip through the Scottish highlands, which is nothing but nature! In the end we found what we were looking for: *Aberlour 10* – rich but delicate.

Lic Zoran Sjanic is as tall as he is fun. And under that surface, lies a great and very helpful source of knowledge and experience. Sometimes when I have one of those questions, Zoran will provide me with all that I ever needed to know. And thank you for showing us Sarajevo.

Lic Hanna Fager is a great friend with whom I have shared many talks, *fikas*, travels and parties during these years. Let us hope we can make it last for a long time.

A few more deserve an honorary mention. Lic Daniel Peterson, my *bother in arms*, is always there to laugh at my *super hacker genius* related issues. Dr Henrik Ohlsson for our *Ella elle l'a* fridays during more than two years as office mates. Lic Sina Khoshfetrat Pakazad serves these meat with a side of meat-barbecues and movies with special effects but no plot. I am surprised and a bit disappointed though that you still have not seen the greatest bad movie of all time, *Frankenfish*, but I guess it is a people problem.

Three years ago Christian, David, Dr Per Skoglar, Dr Peter Bunus, Fredrik and I founded the spin off company SenionLab. It has been a very interesting journey to combine the academic research with the in many ways very different business driven corporate research. This experience has taught me two completely separate ways of thinking.

My will of doing an academic Master thesis project led me into the realm of Dr Juan Nieto and Dr Fabio Ramos, the *dynamic duo* of ACFR. Our great time together made me want to start this second journey towards a PhD. Thank you

for the good times and for the inspiration!

My old friends from the *Y*-years, my childhood and from various corners of the globe, I hope we can maintain our friendships despite the fact that our lives are slowly diverging as time goes by. Hopefully, the completion of this project will provide me with some spare time again that I would be happy to share with you.

I have had the extraordinary luck in life to have been blessed with a family that always supports me. My sister Kajsa and her fiance Mikael have the good sense of also living in Linköping and have created my wonderful little niece Maja. No one can say *mooboo* like you! My parents Gunnar and Eva, thank you for all your never-ending love and support in everything I have ever decided to do. It has made me the one I am today and it is the greatest gift of all.

I would finally like to gratefully acknowledge CADICS, a Linnaeus center funded by the Swedish Research Council, and the Excellence Center at Linköping–Lund in Information Technology, ELLIIT, for providing the funding for this work.

*Linköping, April 2013*
*Jonas Callmer*

# Contents

# Notation

### Abbreviations

| Abbreviation | Meaning |
|---|---|
| CRLB | Cramer-Rao Lower Bound |
| CRM | Corrosion Related Magnetism |
| EKF | Extended Kalman Filter |
| ESDF | Exactly Sparse Delayed-state Filter |
| FIM | Fisher Information Matrix |
| GNSS | Global Navigation Satellite Systems |
| GPS | Global Positioning System |
| HMM | Hidden Markov Model |
| IMM | Interacting Multiple Model |
| IMMU | Inertial Magnetic Measurement Unit |
| INS | Inertial Navigation System |
| KF | Kalman Filter |
| MEMS | Micro-Machined Electromechanical System |
| PDR | Pedestrian Dead Reckoning |
| RADAR | RAdio Detection And Ranging |
| RMSE | Root Mean Square Error |
| SIFT | Scale-Invariant Feature Transform |
| SLAM | Simultaneous Localization And Mapping |
| ZUPT | Zero Velocity Update |

### Estimation

| Notation | Meaning |
|---|---|
| $x(t)$ | State at time $t$ |
| $\dot{x}(t)$ | Derivative of $x(t)$ at time $t$ |
| $x_k$ | State at time step $k$ |
| $y_{1:k}$ | Set of measurement from time step 1 to $k$ |
| $w_k, e_k$ | Process and measurement noise at time step $k$ |
| $Q_k, R_k$ | Process and measurement noise covariance at time step $k$ |
| $T$ | Sampling Time |
| $\hat{x}_{k\|N}$ | State estimate at time step $k$ given measurements up to and including time step $N$ |
| $P_{k\|N}$ | Covariance of state estimate at time step $k$ given measurements up to and including time step $N$ |
| $\mathbf{g}$ | Gravitation vector $(0\ 0\ 9.82)^T$ |
| $\theta$ | Constant parameter vector |
| $\lambda_k$ | Test statistic at time step $k$ |
| $\delta_k$ | Mode at time step $k$ |
| $\mu_k$ | Estimated mode probability at time step $k$ |
| $\psi_k$ | Heading state at time step $k$ |
| $q_k$ | Quaternion orientation at time step $k$ |
| $d_k$ | Deterministic disturbance at time step $k$ |
| $\mathcal{N}(m, \Sigma)$ | Gaussian probability density function with mean $m$ and covariance $\Sigma$ |
| $\chi^2(k, g)$ | Non-central $\chi^2$ distribution with $k$ degrees of freedom and non-centrality parameter $g$ |
| $p(x\|y)$ | Probability density function of $x$ given $y$ |

# Part I

# Background

# 1 Introduction

Localization requires a map and a way to positioning a user in that map. Traditionally the map has been created first. Localization was then solved by placing aiding landmarks in the area like lighthouses for ships, navigation satellites orbiting the planet, or radio beacons on land. These enabled the position to be computed by using distance and/or angle to multiple such landmarks. The methods are called triangulation or multilateration and the position can be acquired by solving an equation system.

This thesis covers the more complex problem of localization in unknown environments. In these scenarios, choosing reliable landmarks in the environment becomes a part of the problem. One must also determine position using a trajectory of earlier positions, where the relationship between different time instances is depending on sensors measuring the system dynamics. The required mathematics is called nonlinear filtering.

The positioning can be performed on any type of unit. Estimating the position of a robot exploring a sewer system, of a car in a city, of a ship in an archipelago or tracking a fire fighter searching through a burning building, all is localization.

If one is outdoors and a Global Navigation Satellite System (GNSS) is available, position estimation becomes straightforward given that the provided measurement accuracy is enough for the application.

There are though many environments where GNSS signals are not available. Such signals are extremely weak making their penetrating ability highly limited. For example indoors, underground or underwater GNSS signals cannot be detected. Even outdoor the GNSS signals can be corrupted. This is commonly caused by the signals being reflected or that the line of sight to a satellite is blocked by trees or

high buildings. Lately, intentional or unintentional jamming of the GNSS signals has emerged as a potential major problem. Jamming the system is very easy since the broadcasted signals are so weak. This makes systems that depend entirely on GNSS quite vulnerable.

Different, redundant means of positioning are therefore required, ones that are tailored for each specific problem. The solutions must be reliable and use all other available information to get the best possible positioning estimate. That is the problem of localization.

## 1.1 Problem Description

Three localization problems have been studied in this thesis. The first is indoor localization for first responders, soldiers and other professional users, the second is surface vessel positioning using a naval RADAR, and the third is underwater sensor positioning using a friendly vessel.

### 1.1.1 Indoor Localization

The problem of indoor localization for professional users has received a lot of attention in the last couple of years Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007); Godha et al. (2006); Woodman and Harle (2009); Grzonka et al. (2010); Aggarwal et al. (2011); Jiménez et al. (2010a); Robertson et al. (2009); Widyawan et al. (2008); Abdulrahim et al. (2011); Jiménez et al. (2010b); Bebek et al. (2010); Angermann and Robertson (2012). Be it firefighters, soldiers or police officers, being able to track the position of each individual user in real time while in a building, is the dream of the operational management. In case something urgent happens, knowing where all the personnel are and where they have been, enables swift and accurate cooperation to solve the problem. Having a positioning system would therefore greatly enhance the safety and efficiency of the personnel.

Figure 1.1 shows the envisioned scenario and positioning presentation such a system will provide in the future. The firefighters are equipped with small lightweight sensors that do not interfere with their ability to do their job. The inertial/magnetic navigation system in the positioning system is supported by collaborative positioning utilizing the distance measured between users, beacons on the trucks providing distance to a fixed position, GPS if available like on the roof and the digital map. The positioning system works in real time, positioning each user with meter level accuracy and broadcasting the information. The information is presented to the operational manager overlaid on an informative 3D map.

For a large venue like a school, hospital, shopping mall or a hotel, the positioning presentation can be put into even more context. In the future, one can envision that not only accurate maps of the building are available but also that smart sensors such as special fire detectors have been installed that can broadcast signals that aid the positioning system, while also providing more detailed information


***Figure 1.1:*** *Visionary illustration of a first responder operation in an urban environment. Firefighters in the building are localized and presented for the commander in real time. Digital maps, collaborative positioning, GPS and beacons on the trucks are used to assist the positioning system. Courtesy of FOI, illustration by Martin Ek.*

about for example the current location and spread of a fire. Also information about where dangerous materials are stored, which hotel rooms that are occupied, in which areas children are likely to be present and so on, can be included to provide an overview of the complete scenario for the management.

The problem studied in this thesis is localization in smaller structures such as residential houses or offices. Larger facilities have the potential of being equipped with designated sensors and maps as described above. For smaller venues such systems are unlikely even in the future, why the positioning system to a large extent has to rely on the sensors brought by the users. The solution should be as simple as possible, using as few sensors as possible and based on as few assumptions about user movements and the environment as possible. The localization system studied in this thesis is therefore a pedestrian dead reckoning system using a foot mounted sensor.

A subproblem of the indoor positioning problem has been given the most attention in this thesis: the problem of heading estimation. Knowing in which direction the user is moving is crucial when determining ones position using dead reckoning. Earlier, heading has not been accurately estimated indoors causing

it to drift and therefore the position estimates to drift. Solutions to the heading estimation problem are presented and incorporated into a position estimation system to enhance the positioning performance.

### 1.1.2 Surface Localization

Modern maritime navigation is highly GNSS centered. It is not only used for positioning but often also as a compass, to track communication satellites and some systems rely on the very accurate measurement of time it produces.

Since GNSS signals are easily jammed, a backup system is needed when navigating in critical environments. Even though pilots are often present in such scenarios, reducing the impact of GNSS failure will further improve the safety of the system.

We present a positioning system based entirely on the measurements from the ship's RAdio Detection And Ranging (RADAR) where the scans are used to estimate the relative position, the velocity and the heading of the vessel.

### 1.1.3 Underwater Localization

A passive surveillance sensor network can position a surface or submerged vessel using underwater sensors. They sense the magnetic field disturbances and acoustic noises caused by the vessel and can thereby determine its position.

One problem is that the exact positions of the sensors are seldom known unless a large amount of time and money have been spent on determining their exact positions. Rapid sensor deployment is therefore difficult since the sensors have to be dropped from a surface vessel and currents can make them move while sinking. Without correct positions of the sensors in the network, accurate tracking of intruding vessels cannot be achieved.

In this thesis we have studied the localization problem of determining the positions of the sensors using a friendly vessel with a known magnetic signature. By knowing where the vessel has been and when, the positions of the sensors can be determined. Now when the true sensor positions are known the network can start undertaking its original task: search for naval intruders.

## 1.2 Contributions

The second part of this thesis constitutes a compilation of five publications.

#### Robust Heading Estimation Indoors

Paper A,

> J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors. *IEEE Transactions on Signal Processing*, 2013a. Submitted.

presents a Kalman filter bank based heading estimation system. Indoors, magnetic heading is not a reliable measurement to use due to frequent and large disturbances. To aid the estimation, measurements of angular velocity from a low grade gyro are incorporated. The Kalman filter bank is used to detect disturbed and undisturbed data segments. To detect filter divergence, a secondary system is used, independent of the filter estimates. The performance of the system is evaluated using more than 500 datasets.

The first author has produced the majority of the ideas, theory and writing and all the implementations, but not all of the data collection.

#### Robust Heading Estimation Indoors using Convex Optimization

Paper B,

> J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors using convex optimization. In *International Conference on Information Fusion*, 2013b. Submitted.

presents a convex optimization based heading estimation system. It states that the gyro signal is correct down to a small gain error and bias. Those two parameters and the initial heading are estimated by tweaking the summed up vector of gyro measurements to match the magnetic heading vector as closely as possible. The matching is done using regularized weighted least squares which can be implemented very cheaply. Also presented is a method to unwrap the magnetic heading vector to enable the matching. The estimation system is shown to work well on more than 500 datasets.

The first author has produced more or less all of the ideas, theory, implementation and writing, but again, not all of the data collection.

#### An Inertial Navigation Framework for Indoor Positioning with Robust Heading

Paper C,

> J. Callmer, D. Törnqvist, and F. Gustafsson. An inertial navigation framework for indoor positioning with robust heading. *IEEE Transactions on Instrumentation and Measurement*, 2013c. Submitted.

presents indoor positioning using a foot mounted inertial magnetic measurement unit. The dead reckoning positioning system is based on a stand still detection system that was in part presented in

> J. Rantakokko, J. Rydell, P. Strömbäck, P. Händel, J. Callmer, D. Törnqvist, F. Gustafsson, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: multisensor systems and cooperative localization. *Wireless Communications, IEEE*, 18(2): 10–18, 2011.

> J. Callmer, D. Törnqvist, and F. Gustafsson. Probabilistic stand still detection using foot mounted IMU. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010b.

To solve the common issue of drift in heading that leads to significant error in position, the system incorporates a robust heading estimation system very similar to the one in Paper B. The positioning performance is shown to improve significantly on two challenging experiments.

The first author has produced the vast majority of the ideas, theory, implementation, experiments and writing in this paper.

#### RADAR SLAM using Visual Features

Paper D,

> J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011.

presents a RADAR based backup system for surface vessel positioning that can be used in case the global navigation satellite system, GNSS, is out. The system estimates relative change in position and heading and velocity using RADAR scans of the surroundings. The RADAR scans are treated like a bird eye's view of the surroundings and consecutive scans are matched using visual features. By studying how the features move over time, the vessel position can be estimated. The system is evaluated using a 32 km experiment.

The first author has produced a majority of the theoretical framework and the writing. Implementation and experiments were produced by Henrik Svensson.

#### Silent Localization of Underwater Sensors using Magnetometers

Paper E,

> J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010a.

presents a positioning system for underwater sensors. An underwater sensor network can be used to track surface or submerged vessels using for example magnetic disturbances and acoustics. For the system to work well, the position of each sensor has to be known which can be hard to determine. This work presents a way to passively position the sensors using a friendly vessel that travel through the area. If the friendly vessel has a known magnetic signature and a known trajectory, the sensor positions can be estimated. The system is evaluated using simulations.

This was joint work between primarily the first and second author who produced the ideas, theory, implementation and most of the writing.

### 1.2.1 Additional Publications

Other publications where the author has contributed that are not covered in this thesis are shortly presented below.

***(a)*** *Skewed image due to rolling shutter and camera turning sideways to the right.*

***(b)*** *Rectified image using sensors of the phone.*

***Figure 1.2:*** *Rectification system for skewed images using the unit sensors. Gyro measurements are used to calculate the distortion of the image which can then be corrected.*

#### Smartphone Stabilization

In

> G. Hanning, N. Forslöw, P.-E. Forssén, E. Ringaby, D. Törnqvist, and J. Callmer. Stabilizing cell phone video using inertial measurement sensors. In *In Proceedings of the IEEE International Workshop on Mobile Vision (IWMV11)*, 2011.

the video stream of an iPhone was stabilized using the inertial sensors. Tracking the orientation of the phone using the gyros and the gravity component, two major errors could be corrected. The first problem is that the image becomes skewed if the phone is moving while the image is taken since the entire image is not recorded at the same time, a so called rolling shutter camera, Figure 1.2. The second problem is that the image becomes unstable if the one holding the camera is moving while filming. Both these problems were solved and the outcome of this master thesis project was the iPhone app *DollyCam*. The author served as supervisor to Nicklas Forslöw during the master thesis project that was the foundation of the paper. It was awarded best paper at the workshop.

#### Vehicle Tracking using Magnetometers

The two papers

> N. Wahlström, J. Callmer, and F. Gustafsson. Single target tracking using vector magnetometers. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.

> N. Wahlström, J. Callmer, and F. Gustafsson. Magnetometers for tracking metallic targets. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

are about vehicle tracking using a three axis magnetometer. Passing vehicles disturb the earth magnetic field, making it possible to estimate the position and direction of each passing vehicle. Multiple filters were initiated from all possible directions once a vehicle was detected. The probability of each trajectory was estimated and only one estimate survived. The results from an experiment is shown in Figure 1.3. The publications are based on a master thesis project undertaken by Niklas Wahlström that was supervised by the author.

#### Geo-referencing for UAV Positioning

As backup for GNSS,

> F. Lindsten, J. Callmer, H. Ohlsson, D. Törnqvist, T. B. Schön, and F. Gustafsson. Geo-referencing for UAV navigation using environmental classification. In *Proceedings of 2010 International Conference on Robotics and Automation (ICRA)*, 2010.

covered the problem of using preexisting maps and environmental classification to create a measurement of the global position of an Unmanned Aerial Vehicle (UAV).

Photos from a downwards facing camera on the UAV were classified into grass, houses, roads etc, which could be matched to a map of the area, Figure 1.4. Using a rotation invariant probabilistic class matching system, dubbed 'donuts', a likelihood for the position of the vehicle could be provided. Merged with a inertial/visual odometry based positioning system, the overall drift in the position estimate could be significantly reduced. The main contribution of the author was within the image classification.

*(a) The vehicle is coming from the rear turning right.*

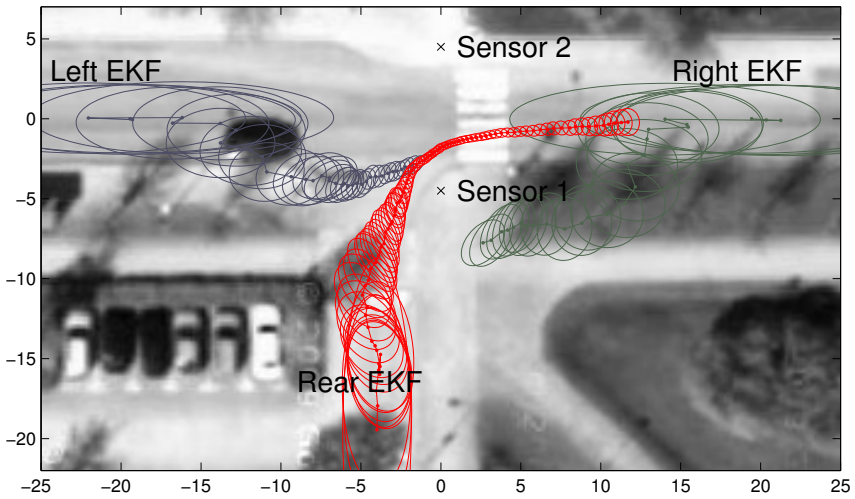


*(b) The trajectories according to three EKFs with different vehicle position initializations.*

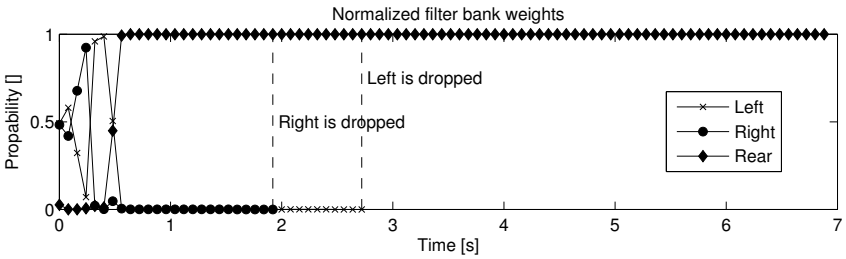


*(c) The probabilities that the vehicle is coming from the left, the right and the rear. Only the hypothesis that the vehicle is coming from the rear survives and the other two are dropped.*

***Figure 1.3:*** *Tracking experiment result with three differently initialized extended Kalman filters estimating the trajectory of the vehicle in Figure 1.3a.*

*(a)* Image from camera onboard a UAV.

*(b)* Extracted superpixels.

*(c)* Superpixels classified as grass, asphalt or house.

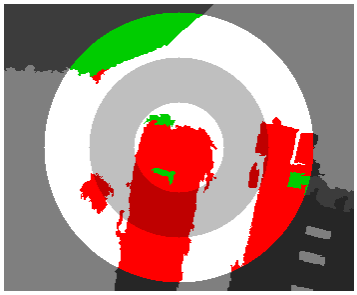

*(d)* Three circular regions used for computing class histograms.

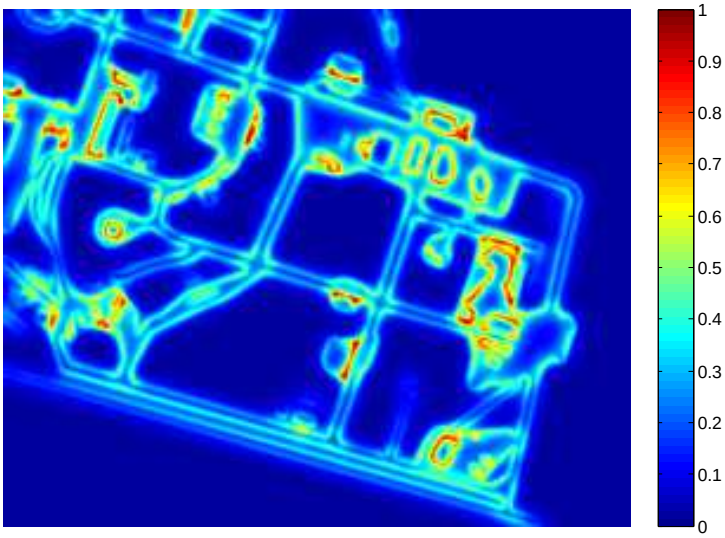


*(e)* Calculated likelihood for all available locations.

***Figure 1.4:*** *From raw image to position likelihood. Image is classified and class histograms are used to calculate position likelihood.*

### Loop Closure Detection for SLAM

The last two publications are spinoffs from the master thesis project undertaken by Karl Granström and the author at Australian Centre for Field Robotics, Sydney University in 2007/08.

> K. Granström, J. Callmer, F. Ramos, and J. Nieto. Learning to detect loop closure from range data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.
>
> J. Callmer, K. Granström, J. Nieto, and F. Ramos. Tree of words for visual loop closure detection in urban SLAM. In *Proceedings of the 2008 Australasian Conference on Robotics and Automation (ACRA)*, 2008.

Both papers are about loop closure detection methods for large scale urban simultaneous localization and mapping (SLAM). Granström et al. (2009) we presented a laser scan based matching method that matched 360° 2D range slices of the surroundings using rotation invariant features. Adaboost were used to produce a reliable scan matcher. Callmer et al. (2008) detected loop closures using photos of the surroundings. Visual features were extracted from the images and approximated as one of a large number of predefined words. Two images were matched by comparing the list of words they contained. Both methods were incorporated into a SLAM estimation system to produce maps of the area, Figure 1.5.

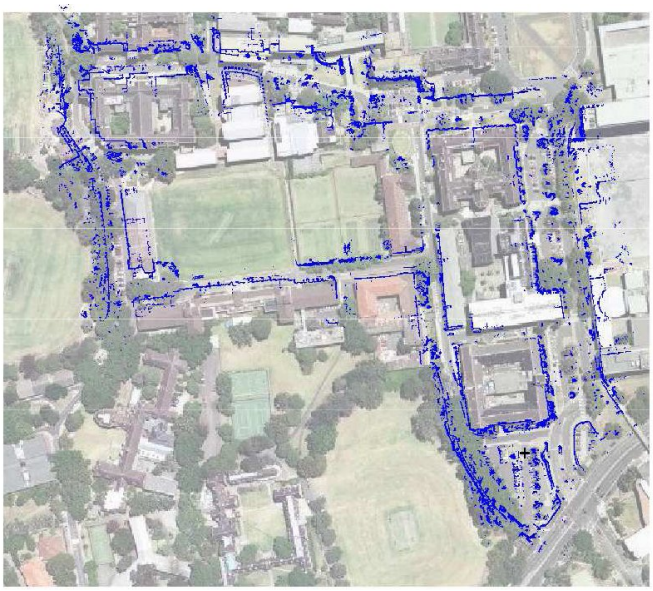

***Figure 1.5:*** *Laser map of based on the results of a SLAM experiment, overlaid on an aerial photograph.*

### 1.2.2 SenionLab

During the fall of 2010, the spinoff company SenionLab was cofounded by the author. The company provides indoor navigation solutions for primarily cell phones, Figure 1.6. The positioning system is intended for facilities like shopping malls, airports, hospitals, stadiums etc. At the time of writing, the largest commercial deployment is 50 malls in Singapore.

***Figure 1.6:*** *Indoor positioning using smartphone. Courtesy of SenionLab AB.*

A patent application on the technology has been filed.

> C. Lundquist, P. Skoglar, F. Gustafsson, D. Törnqvist, and J. Callmer. Method and device for indoor positioning. US Patent Application 20120203453, August 8, 2012.

The system uses all available sensors in the device to position the user. Radio environment mapping and sensing enables accurate positioning that is fused with a dead reckoning system for a smooth user experience.

Several challenges have been met that greatly illustrate the difference between commercial product research and development on one side and academic research on the other. A major challenge is for example how to handle many different types of sensors with different sampling times to make the system work as well as possible on all types of known and even unknown units. Even more challenging is the difference in what different sensors measure. The received signal strength indicator of WiFi measurements differ for example greatly between different phones.

Besides sensor variabilities, the application scenario is a tremendous source of

problems that need to be solved. How different users handle the unit and how they move is one such. Even greater are practical matters like errors in floor plans, reconstructions, the frequent magnetic disturbances, WiFi access points that move and WiFi access points that change broadcasted signal strength depending on the current workload. Also, to produce maps of the radio environment, one must have measurements of it, but one must also know exactly where each measurement was taken. And in order to make deployment cheap and fast, a logging tool must be created that is fault tolerant and intuitive enough to be handled by just about anybody. The logging tool should preferably also signal to the one doing the logging that what he or she is doing is important and must be done meticulously. These are challenging and important issues one does not normally experience when doing academic research.

But the practical issues can also lead to new research ideas. In fact, the experience of magnetic disturbances in practice has been an inspirational source for Papers A and B.

In the end, commercial deployment of a solution to a research problem introduces whole new problems that are at least as hard as the original problem. Constructing a system that should work a billion times on a million different units is very different from academic research.

## 1.3 Thesis Outline

This thesis is divided into two parts. The first part provides background theory and gives context to the second part which is constituted of the five edited publications.

Chapter 2 is a brief introduction to the basic sensor fusion tools of sensors, modeling and estimation theory that is the fundament of all publications. It ends with a discussion about a very central problem in this thesis: estimation using disturbed measurements. Chapter 3 describes the problem of indoor positioning. Indoor navigation using dead reckoning for first responders is discussed and one of the systems later used in Paper C is described in more detail. Chapter 4 summarizes the first part of the thesis with conclusions and a discussion about future work.

# 2
# Sensor Fusion

Sensor fusion is the problem of estimating some properties $x_t$ of a unit, using sensors that provide measurements $y_t$ that depend on $x_t$. In order to do this, models of how $y_t$ is related to $x_t$ and of how $x_t$ changes over time, are used. The former are called measurement models and the latter are called process models.

The properties $x_t$ are called states and can represent any sought system property. The states can for example be related to the sensor platform representing the position or orientation of the unit, they can be some unknown constant properties such as the unit weight or they can be of the surrounding environment such as the positions of environmental landmarks. In the problem of localization the states are commonly position, velocity and orientation of the unit which are key features representing 'where is the unit?' and 'where is the unit going?'.

The states $x_t$ are estimated using a filter that fuses the information from all the sensors and the models. Besides from the state estimates $\hat{x}_t$, the filter also provides an estimate of the uncertainties of $\hat{x}_t$.

The joint estimate $\hat{x}_t$ is in some sense better than what one could get using the sensors individually. The meaning of better is application dependent and could for example mean more accurate estimates, more robust estimates or that the same estimation precision can be achieved using fewer sensors.

A schematic overview of a sensor fusion framework is given in Figure 2.1. The sensor measurements enter the estimation system that relates these to the system states using the measurement models. The states are updated using the process models describing the system, which are then fused with the measurements. The system output is the state estimates which can be used by other applications.

Three components are needed in sensor fusion: sensors producing measurements

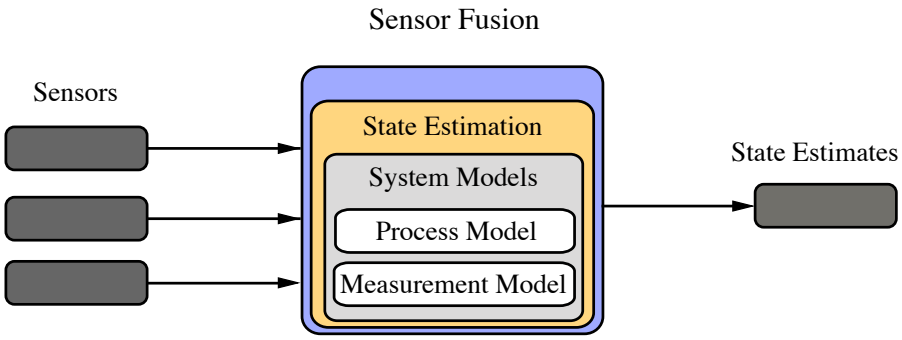


***Figure 2.1:*** *Overview of a sensor fusion framework. Sensor data is fused with dynamic system models using measurement models to produce state estimates.*

$y_t$ that are related to the system states $x_t$, models that describe the dynamic properties of the system and the measurements, and a state estimation system that produces the state estimates. This chapter will describe all three parts and provide the background theory on the subject needed for the publications part of the thesis.

## 2.1 Sensors

The sensors produce measurements that relate to the sought system properties in some way. Which sensors that are used is highly application dependent.

In localization problems, the sensors are usually of one of two kinds. The first option is that they produce measurement that are related to the unit directly such as its movement, position, direction etc. The second option is that they are indirectly related to the unit, measuring for example some aspect of the surrounding environment that will change when the unit properties change. They can be cameras filming the surroundings, range finders measuring the distances to objects around the unit, or magnetometers measuring the surrounding magnetic field, among others.

The sensor primarily used in this thesis is an inertial magnetic measurement unit (IMMU). It is actually not just a sensor but a sensor unit containing nine separate sensors: three accelerometers, three gyros and three magnetometers. It has been used in Papers A, B and C. In Paper D, a RADAR sensor was used for maritime localization experiments. Paper E contains only simulations why the sensor data came from a simulated three axis magnetometer and a pressure sensor. The IMMU and the RADAR sensor are described in this section and also a short description of GNSS.

***Figure 2.2:*** *An Xsens MT motion sensor, courtesy of Xsens Technologies B.V.*

### 2.1.1 Inertial Magnetic Measurement Unit

An inertial magnetic measurement unit contains an accelerometer, a gyroscope and a magnetometer, all three dimensional. The accelerometer measures acceleration, the gyro measures the angular velocity and the magnetometer measures the magnetic field. Besides from these, the IMMU often contains a thermometer to enable correcting temperature related sensor errors.

There are many different kinds of IMMUs with different price, size and precision but in this section we will focus on only one type: micro-machined electromechanical systems (MEMS). MEMS sensors are small, rugged, low cost, lightweight and low on power consumption, making them popular to include in all sorts of devices. On the downside, their performance is in many aspects quite poor even though it is constantly improving. A MEMS IMMU can be made to work better through calibration, but this is labor intensive which is expensive.

A MEMS IMMU is also a strap-down system, meaning that the sensors are mounted on the device making the measurements in the body frame.

The IMMU used in the standstill detection in Chapter 3 is an Xsens MT motion sensor, Figure 2.2. The IMMU used in Paper C was an MicroStrain 3DM-GX3-25. The signals were sampled in 100 Hz using a 16 bit A/D converter.

#### Accelerometer

The accelerometer actually measures the specific force which is a type of acceleration. Specific force is defined as non-gravitational force per unit mass, meaning it is the acceleration relative to free-fall. A free-falling accelerometer therefore experiences no specific force while a sensor at rest senses the normal force from the surface that cancels the gravity. An accelerometer at rest therefore measures the gravitational constant $g$ but pointing upwards, not downwards.

There are mainly two types of accelerometers: mechanical sensors and solid state sensors. The mechanical sensor measures how a suspended mass is displaced due to an applied force. Using Newtons second law $F = ma$, the acceleration can be

measured. A solid state sensor is for example the surface acoustic wave (SAW) accelerometer. It uses a mass attached to a beam that is vibrating at a particular frequency. When a force is applied, the beam bends, changing the frequency. Thereby the force can be measured. A MEMS accelerometer can be based on any of these techniques.

The main errors associated with mems accelerometers, besides the ever present additive white noise, are

- Bias - the sensor value has a slight offset
- Temperature - a change in temperature gives a change in measured output
- Calibration errors - such as alignment errors, gain errors and so on

Bias errors can sometimes be estimated depending on the application and other sensors available. Temperature errors are commonly handled by the sensor unit. But since the temperature errors are often highly nonlinear the effect is often not completely removed. Calibration errors are very hard to estimate if present, especially alignment errors.

### Gyro

The gyro measures angular velocity, i.e. rate of turn. The gyro is often the weakest point in an inertial navigation system, INS. Accurately estimating the orientation of the device is crucial in such a system and the orientation is tracked using the gyro.

A high grade gyro is often based on optics such as ring laser gyros or fibre optic gyros. The sensor is based on light inference. Two light beams are shone into opposite ways of a track. The track can be an optical fibre or a mirror path, for example. When the beams return, the inference, i.e. phase shift, reveals if one of the beams has travelled a shorter path due to that the path has rotated. This is called the Sagnac effect. Optical gyros often have high precision, but are hard to reduce too much in size since a shorter path means worse precision.

A MEMS gyro is based on that a moving element that is effected by a rotation gives away a force in the perpendicular direction

$$F^c = -2m(\omega \times v). \tag{2.1}$$

This is known as the Coriolis effect. The velocity is often represented using a vibrating mass to create $F^c$. Since $v$ is a vibration, $F^c$ changes direction with the direction change in $v$, why it is also vibrating. Today the MEMS gyros cannot match the precision of the optical gyros.

As for the MEMS accelerometers, the main sources of error are

- Bias
- Bias stability - the bias actually moves around slightly and is not as constant as a bias should be

- Temperature effects
- Calibration errors.

and the additive white noise. Since the rotations are often integrated over time to produce an orientation estimate, the fact that the bias does not even have the decency to be still is integrated into a significant orientation drift over time.

#### Magnetometer

Magnetometers measure the magnetic fields. The magnetic fields consists of the earth magnetic field and local magnetic disturbances. If the earth magnetic field is stronger than the disturbances, information of the direction of magnetic north is available.

There are many different approaches to magnetic sensing such as Hall effect sensor, magneto-diode, magneto-transistor etc. MEMS magnetometers are often based on the Lorentz-force which acts on a current-carrying conductor in a magnetic field.

$$\mathbf{F} = q[\mathbf{E} + (\mathbf{v} \times \mathbf{B})] \tag{2.2}$$

where $q$ is the charge, $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic field and $\mathbf{v}$ is the velocity of the charge.

This force can be measured in different ways, for example by sensing the strain this force applies on piezo-resistors or by sensing a frequency shift in a beam caused by the force. One can also detect the force by studying the displacement of the MEMS structure.

### 2.1.2 RADAR

A pulse RADAR sends out radio waves in different directions which are reflected or scattered when hitting an object. The reflected signals are picked up by a receiver, usually at the same place as the transmitter, and the time of flight for the signal is calculated. This time is proportional to the distance to the object that reflected the signal and the heading of the sensor when the signal was transmitted gives the direction to the object. The strength of the reflected signal can also provide some information about the properties of the reflecting object.

Two measurements are provided by each reflected wave: range and angle from the sensor to the object.

$$r = \sqrt{(s_x - p_x)^2 + (s_y - p_y)^2} \tag{2.3}$$

$$\alpha = \arctan \frac{s_y - p_y}{s_x - p_x} \tag{2.4}$$

where $s = (s_x, s_y)$ is the position of the reflecting structure and $p = (p_x, p_y)$ is the position of the RADAR sensor. Since the uncertainties in angle and range are independent, the total measurement uncertainties will be banana shaped.

A naval RADAR commonly rotates with a constant speed, transmitting and receiving in one direction at a time. The reflections are plotted in the current direction when they are received. This gives a circular image of the surrounding islands and vessels that is updated one degree at a time. By saving one 360° RADAR sweep as an image, a view of the surroundings is provided.

The RADAR sensor used in Paper D was a military one making the characteristics of that sensor secret. What we do know is that it had a range of roughly 5 km and a range resolution of about 5 meters. It rotates one revolution in 1.5 seconds giving measurements in roughly 2000 directions.

One way of using a RADAR in localization is to take some strong reflections in a full 360° RADAR scan and try to detect them again in the next scan. The objects creating these reflections are called landmarks and are assumed stationary. By measuring the distance and heading to the landmarks and see how these change over time, how the RADAR equipped unit is moving can be estimated. If some landmarks move in a manner that is inconsistent with the other landmarks, it is probably a different unit and the reflections should not be used for localization.

### 2.1.3 Global Navigation Satellite System

Global Navigation Satellite System use multiple satellites and triangulation to determine the position of a user anywhere on earth. The most well known such system, the Global Positioning System (GPS), provides a positioning accuracy of about 10 meters. Besides from location, GPS also gives very accurate estimates of current time, making it useful also in applications where only accurate time and not position is needed. This is for example used in cellphone base station synchronization for some systems. The system consists of 30 satellites and free line of sight to at least 4 of them is required for the positioning to work.

Other systems exist or are planned. The Russian GLONASS system mostly covers the northern hemisphere, in particular Russia, and is today short of the 24 satellites needed to cover the whole planet. The European Galileo system will use 30 satellites to cover the entire planet and the full deployment is expected to be finished in 2019. Also a Chinese system, COMPASS, using 35 satellites to cover the planet will be deployed in the future. As of today, a smaller system covering only China and the immediate surroundings is in place. A future GNSS receiver, using signals from all systems will pretty much always have free line of sight to at least 4 satellites. This will give accurate positioning also in places that are difficult to cover today such as urban canyons.

One shortcoming with GNSS systems is the weakness of the signals. The signal is weaker than the background noise and only because the receivers know what to look for can the signals be found. This makes the system sensitive to signal disturbances due to intentional or unintentional jamming. Today, GPS jammers that can easily knock out all GPS reception in an area of many square kilometers are available at a low cost; Economist (2011); Grant et al. (2009). This problem and a suggested solution for maritime vessels is discussed in more detail in Paper D.

## 2.2 Models

In estimation, mathematical models are used to describe how the states are related to eachother and to the measurements.

A process model describe the dynamic properties of the system by stating how the states depend on one-another and on additional inputs. For a vehicle model, the process model describes how the velocity states translate into a change in position states over time for example. It will also put restrictions on a system by stating that a vehicle cannot not travel sideways for example. Dynamic system models are often relating the states to eachother using differential equations. To simplify implementation, these models are most often approximated as discrete time difference models.

A measurement model relates what is measured by the sensors to the unit states. The measurements can be of the states themselves or they can be functions of one or more states. The mathematical models used to describe the relationship between the measurements and the states are often nonlinear functions.

### 2.2.1 Continuous Models

The models are commonly on a state space form where a state vector $x(t)$ describes the system properties at time $t$. The process model is $f(\cdot)$ and the measurement model is $h(\cdot)$.

The fundamental continuous time model is

$$\dot{x}(t) = f(x(t), u(t), w(t)) \tag{2.5}$$

$$y(t) = h(x(t), u(t), e(t)) \tag{2.6}$$

where $u(t)$ is a known input signal and $w$ and $e$ are model and measurement noise terms, respectively. $f(\cdot)$ and $h(\cdot)$ are in general nonlinear functions.

Even though the process model (2.5) is often based on fundamental relationships between states described by differential equations, some simplifications have always been made of the true system. The dynamic model is therefore associated with a process noise, which is the assumed input that is driving the true system. The process noise should also incorporate the model uncertainties.

Related to each measurement in the measurement model (2.6), is a measurement noise $e(t)$. No matter the sensor, there is always a noise present in the measurements. The noise term $e(t)$ therefore reflects the quality of the sensor, with larger noise covariance terms for poor sensors. For a presentation on random signals in continuous time, see e.g. Jazwinski (1970).

### 2.2.2 Discrete Time Models

Estimation methods are primarily based on discrete time systems due to the implementational simplifications that follow. Therefore, continuous time models need to be discretized before they can be used in an estimation systems.

### Discretization

Discretization means that the differential equations in a continuous time model are replaces by approximate difference models that resemble the original ones.

Most often discretization is a complex task. One exception is a linear continuous time system.

$$\begin{aligned} \dot{x}(t) &= Fx(t) + Gw(t) \\ y(t) &= Cx(t) + e(t) \end{aligned} \tag{2.7}$$

To discretize such a system, one must first assume that $w_k$ is piecewise constant over the sampling interval $T$. The matrices of the sampled systems can then be computed as

$$A = e^{FT} \tag{2.8}$$

$$B = \int_0^T e^{F\tau} d\tau G \tag{2.9}$$

giving the discrete time linear system

$$\begin{aligned} x_{k+1} &= Ax_k + Bw_k \\ y_k &= Cx_k + De_k. \end{aligned} \tag{2.10}$$

where $w_k \sim \mathcal{N}(0, Q_k)$ and $e_k \sim \mathcal{N}(0, R_k)$ Note that the model matrices $A$, $B$, $C$ and $D$ need not be constant.

For most other cases sampling a continuous time system is quite challenging. For details see Gustafsson (2010).

### General Discrete Time model

A general description of a physical system as a state space model in discrete time is

$$\begin{aligned} x_{k+1} &= f(x_k, u_k, w_k) \\ y_k &= h(x_k, e_k) \end{aligned} \tag{2.11}$$

An important special case is when the process and measurement noises are modeled as additive

$$\begin{aligned} x_{k+1} &= f(x_k, u_k) + w_k \\ y_k &= h(x_k) + e_k. \end{aligned} \tag{2.12}$$

It is an intuitively straightforward model with a deterministic part utilizing basic physical properties and a random part representing everything that is unknown that affects the system. An example of a system with partially nonlinear dynamics and measurements is given in Example 2.1.

Often the input $u_k$ is unknown making the model

$$\begin{aligned} x_{k+1} &= f(x_k) + w_k \\ y_k &= h(x_k) + e_k \end{aligned} \tag{2.13}$$

This is the case when one wants to estimate the properties of a system one does not control, like in target tracking.

The linear system (2.10) is a very common special case of modeling since it allows Kalman filter theory to be applied when solving the problem, if the process and measurement noises are assumed Gaussian.

**2.1 Example**

Model of an inertial navigation system estimating the sensor position using an accelerometer and a gyro and also global measurements of position from GPS.

The states position $p$, velocity $v$ and acceleration $a$, all in global coordinates. Orientation is represented by quaternions $q$ and angular velocity $\omega$ in the local coordinate system. The measurements acceleration $y_a$ and angular velocity $y_\omega$, both measured in the local coordinate system and global position $y_p$ from the GPS receiver. All measurements have additive noise Gaussian zero mean noise $e$.

To reduce the model complexity, the acceleration state $a$ can be replaced by the acceleration measurements $y_a$ in the dynamic model. The corresponding simplification can be done to the orientation, replacing $\omega$ by $y_\omega$ in the dynamic model. The main difference this induces is that high frequency components in $y_a$ and $y_\omega$ are not filtered out, Callmer (2011).

Quaternion dynamics and properties are described in Appendix A but a very brief explanation will be given here. $S'(q_k)\omega_k$ describes how local angular velocities translate into changes in quaternions. $R(q_k)$ is the rotation matrix from the global to the local coordinate system which is based on the quaternions $q_k$.

$$\begin{pmatrix} p_{k+1} \\ v_{k+1} \\ q_{k+1} \end{pmatrix} = \begin{pmatrix} I & TI & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \begin{pmatrix} p_k \\ v_k \\ q_k \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2} I & 0 \\ TI & 0 \\ 0 & \frac{T}{2} S'(q_k) \end{pmatrix} \begin{pmatrix} R^T(q_k) y_{a,k} - \mathbf{g} + e_{a,k} \\ y_{\omega,k} + e_{\omega,k} \end{pmatrix} \tag{2.14}$$

$$y_{p,k} = p_k + e_{p,k} \tag{2.15}$$

where $\mathbf{g}$ is the gravity component.

### 2.2.3 Dynamic Model Restructuring

In some cases we have a dynamic model like (2.12) where the input $u_k$ is known. We then have

$$x_{k+1} = f(x_k, u_k) + w_k \tag{2.16}$$

$$x_k = f(x_{k-1}, u_{k-1}) + w_{k-1} \tag{2.17}$$

which can be rewritten as

$$x_{k+1} = f(f(x_{k-1}, u_{k-1}) + w_{k-1}, u_k) + w_k. \tag{2.18}$$

This procedure can be repeated over and over until $x_{k+1}$ is only depending on the initial states $x_0$, the known inputs $u_{0:k}$ and the process noise terms $w_{0:k}$. In the general case the full expression becomes hideous and unmanageable, but in the special case

$$f(x_k, u_k) = x_k + u_k \tag{2.19}$$

(2.18) becomes

$$x_{k+1} = x_0 + \sum_{i=0}^{k} u_i + \sum_{i=0}^{k} w_i. \tag{2.20}$$

That means that in those cases, $x_{k+1}$ can be written as a simple function of only initial states, inputs and noise terms.

The resulting measurement model is

$$\begin{aligned} y_k &= h(x_k) + v_k \\ &= h(x_0 + \sum_{i=0}^{k-1} u_i + \sum_{i=0}^{k-1} w_i) + v_k. \end{aligned} \tag{2.21}$$

In this thesis, this form has been applied to the one dimensional heading estimation problem. The heading state $\psi_k$ is driven using the gyro measurements $y_k^\omega$ which have gaussian measurement noise $e_k$.

$$\psi_{k+1} = \psi_k + T y_k^\omega + T e_k \tag{2.22}$$

can therefore be written as

$$\psi_{k+1} = \psi_0 + T \sum_{i=0}^{k} y_i^\omega + T \sum_{i=0}^{k} e_i. \tag{2.23}$$

The state vector $\psi_{0:k+1}$ can hence be written on batch form very easily. Such an approach also means that we assume that the gyro sensor has the actual bandwidth to capture all the dynamics of the system.

This model structure will be used in Papers B and C where heading is estimated in batch form.

## 2.3 Estimation Theory

The estimation problem is the problem of estimating the posterior distribution of the states given the measurements, $p(x_k|y_{1:k})$. The states are often intricately related to the measurements, making them difficult to estimate. With the use of Bayes' theorem

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \tag{2.24}$$

the problem can be reformulated into three straightforward parts. $p(y|x)$ is the likelihood of receiving the measurement $y$ given the states $x$, $p(x)$ is the prior probability of the states incorporating all our previous knowledge about the states and $p(y)$, the probability of the measurement, normalizes the state probabilities.

Estimation theory can be divided into two cases: linear and nonlinear estimation. Linear estimation is straightforward and the results are trustworthy but the problem is in reality rare. Nonlinear estimation is more difficult and the solutions are prone to diverge but unfortunately the problem is very common.

In this section we will first present the Kalman filter used for linear estimation problems and the extended Kalman filter used for slightly nonlinear problems. We will then present Kalman filter banks which is a type of filters where multiple Kalman filters with different system assumptions, so called modes, are run in parallel. The final estimate is a mixture between all the filter estimates. This is followed by a presentation of Hidden Markov Models which can be used when only mode is sought. The section ends with a discussion about estimation using disturbed measurements and a suggested solution outline.

### 2.3.1 Kalman Filter

The linear estimation problem where a discrete time linear system (2.10) is assumed to have Gaussian process and measurement noise, is optimally solved using the Kalman filter, Kalman (1960). Since the problem is Gaussian, estimating the mean and the covariance of $x_k$ provides the entire solution to $p(x_k|y_{1:k})$.

The Kalman filter works in a two step procedure with a time update and a measurement update. The time update predicts the future states $\hat{x}_{k+1|k}$ using the process model. Since the model is not perfect, the process noise covariance $Q_k$ is added to the state covariance $P_{k|k}$ to illustrate the increase in estimate uncertainty that the model introduces.

The measurement update uses the difference between the measurement and the predicted measurement, the so called innovation, to update the states. How much the new measurement should affect the states is decided by the Kalman gain, $K_k$. $K_k$ depends on $Q_k$ and the measurement noise covariance $R_k$ which describes how trustworthy the measurements are. The relation between these two parameters determines the filter performance. If $Q_k$ is small in relation to $R_k$, the model is deemed more reliable than the measurements and vice versa. The ratio between $Q_k$ and $R_k$ affects the state estimates while the magnitudes of $Q_k$ and $R_k$ determines the size of the state estimate covariances.

The equations defining the Kalman filter are shown in Algorithm 1.

### 2.3.2 Extended Kalman Filter

If the problem is of the form (2.11) or (2.12) and only mildly nonlinear, the extended Kalman filter (EKF) can be applied. It approximates the nonlinearities using a first order Taylor approximation around the latest state estimate and then

**Algorithm 1** Kalman Filter

**Require:** Signal model (2.10), initial state estimate $\hat{x}_{0|0}$ and covariance $P_{0|0}$.

1: Time Update

$$\begin{aligned} \hat{x}_{k+1|k} &= A\hat{x}_{k|k} \\ P_{k+1|k} &= AP_{k|k}A^T + Q \end{aligned} \tag{2.25}$$

2: Measurement Update

$$\begin{aligned} K_{k+1} &= P_{k+1|k}C^T\left(CP_{k+1|k}C^T + R\right)^{-1} \\ \hat{x}_{k+1|k+1} &= \hat{x}_{k+1|k} + K_{k+1}\left(y_{k+1} - C\hat{x}_{k+1|k}\right) \\ P_{k+1|k+1} &= P_{k+1|k} - K_{k+1}CP_{k+1|k} \end{aligned} \tag{2.26}$$

applies the Kalman filter to the linearized problem. Convergence cannot be guaranteed, in particular because the EKF gives the optimal solution to the wrong problem. That means, the solution to the linearized problem is the optimal one, unfortunately the linearized problem is not the true problem. Despite this, the EKF most often works well in practice. The EKF is primarily used in Papers C and E.

Starting with the nonlinear function (2.12), a first order Taylor expansion of the measurement function $h(\cdot)$ around the linearization point $\hat{x}_k$ is

$$h(x_k) \approx h(\hat{x}_k) + h'_x(\hat{x}_k)(x_k - \hat{x}_k) \tag{2.27}$$

where $h'_x(\hat{x}_k)$ is the Jacobian

$$h'_x(\hat{x}_k) = \left.\frac{\partial h(x_k)}{\partial x_k}\right|_{x_k=\hat{x}_k} \tag{2.28}$$

The measurement model can now be approximated according to

$$\begin{aligned} y_k &= h(\hat{x}_k) + h'_x(\hat{x}_k)(x_k - \hat{x}_k) + e_k \\ y_k - h(\hat{x}_k) + h'_x(\hat{x}_k)\hat{x}_k &= h'_x(\hat{x}_k)x_k + e_k \\ \bar{y}_k &= h'_x(\hat{x}_k)x_k + e_k. \end{aligned} \tag{2.29}$$

Correspondingly the dynamic model can be expanded around $\hat{x}_k$ as

$$f(x_k) \approx f(\hat{x}_k) + f'_x(\hat{x}_k)(x_k - \hat{x}_k) \tag{2.30}$$

giving the new model

$$\begin{aligned} x_{k+1} &= f(\hat{x}_k) + f'_x(\hat{x}_k)(x_k - \hat{x}_k) + w_k \\ x_{k+1} - f(\hat{x}_k) + f'_x(\hat{x}_k)\hat{x}_k &= f'_x(\hat{x}_k)x_k + w_k \\ \bar{x}_{k+1} &= f'_x(\hat{x}_k)x_k + w_k \end{aligned} \tag{2.31}$$

where

$$f_x'(\hat{x}_k) = \left.\frac{\partial f(x_k)}{\partial x_k}\right|_{x_k=\hat{x}_k} \tag{2.32}$$

If the signal model is (2.11), the Jacobians with respect to the noise parameters $w$ and $e$ are also needed.

$$f_w'(\hat{x}_k) = \left.\frac{\partial f(x_k, w_k)}{\partial w_k}\right|_{x_k=\hat{x}_k} \qquad h_e'(\hat{x}_k) = \left.\frac{\partial h(x_k, e_k)}{\partial e_k}\right|_{x_k=\hat{x}_k} \tag{2.33}$$

The extended Kalman filter is summarized in Algorithm 2.

**Algorithm 2** Extended Kalman Filter

**Require:** Signal model (2.11), initial state estimate $\hat{x}_{0|0}$ and covariance $P_{0|0}$.

1: Time Update

$$\begin{aligned} \hat{x}_{k+1|k} &= f(\hat{x}_{k|k}) \\ P_{k+1|k} &= f_x'(\hat{x}_{k|k})P_{k|k}f_x'(\hat{x}_{k|k})^T + f_w'(\hat{x}_{k|k})Qf_w'(\hat{x}_{k|k})^T \end{aligned} \tag{2.34}$$

2: Measurement Update

$$\begin{aligned} S_{k+1} &= h_x'(\hat{x}_{k+1|k})P_{k+1|k}h_x'(\hat{x}_{k+1|k})^T + h_e'(\hat{x}_{k+1|k})Rh_e'(\hat{x}_{k+1|k})^T \\ K_{k+1} &= P_{k+1|k}h_x'(\hat{x}_{k+1|k})^T S_{k+1}^{-1} \\ \hat{x}_{k+1|k+1} &= \hat{x}_{k+1|k} + K_{k+1}\big(y_{k+1} - h(\hat{x}_{k+1|k})\big) \\ P_{k+1|k+1} &= P_{k+1|k} - P_{k+1|k}h_x'(\hat{x}_{k+1|k})^T S_{k+1}^{-1} h_x'(\hat{x}_{k+1|k})P_{k+1|k} \end{aligned} \tag{2.35}$$

### 2.3.3 Optimization Formulation

A system with state vector $x_k$ with dynamic model

$$x_{k+1} = Ax_k + Bu_k + w_k \tag{2.36}$$

and measurement model

$$y_k = Cx_k + e_k \tag{2.37}$$

can, as was described in Section 2.3.1, be estimated using a Kalman filter.

The Kalman filter measurement update can be written as a one step optimization problem.

$$\begin{aligned} \hat{x}_k = \arg\min_{x_k} \quad & (y_k - Cx_k)^T R_k^{-1}(y_k - Cx_k) + \\ & (x_k - Ax_{k-1} - Bu_k)^T P_{k|k-1}^{-1}(x_k - Ax_{k-1} - Bu_k) \end{aligned} \tag{2.38}$$

where $x_{k-1}$ is the previous estimate that is fixed and $P_{k|k-1}$ is the time updated state covariance. This is an unconstrained regularized weighted least squares problem. The closed form solution is the Kalman filter.

To instead produce a smoothing estimate $\hat{x}_{0:k}$, (2.38) can be written on batch form.

Thereby, the entire vector $\hat{x}_{0:k}$ is produced in one step using all measurements $y_{0:k}$. The batch formulation includes both time and measurement update and also a prior on the initial state $x_0$.

$$\begin{aligned}\hat{x}_{0:k} = \arg\min_{x_{0:k}} \quad & (y_{0:k} - Cx_{0:k})^T R_{0:k}^{-1}(y_{0:k} - Cx_{0:k}) + \\ & (\bar{A}_k x_{0:k} - Bu_{1:k})^T Q_{1:k}^{-1}(\bar{A}_k x_{0:k} - Bu_{1:k}) + x_0^T P_0^{-1} x_0\end{aligned} \tag{2.39}$$

where $\bar{A}_k$ is

$$\bar{A}_k = \begin{pmatrix} I & -A & 0 & \dots & 0 \\ 0 & I & -A & \dots & 0 \\ \vdots & & \ddots & \ddots & \\ 0 & 0 & \dots & I & -A \end{pmatrix} \tag{2.40}$$

In most practical applications, estimating $\hat{x}_{0:k}$ using (2.39) is unnecessary. If one wants to include future measurements in an estimate, a fixed-lag smoothing is computationally more efficient and the estimate is very similar to the one produced using (2.39) due to the exponential forgetting factor of old measurements. More on smoothing in Gustafsson (2010).

However, in some application this formulation can be very useful. In this thesis we have one such scenario, in Papers B and C, where heading is estimated on batch form to handle large and frequent disturbances in the measurements.

### 2.3.4 Kalman Filter Banks

In some cases one has a model that depend on a discrete mode $\delta_k$. It can represent many different features such as type of dynamic model (turning or straight) in target tracking, a faulty or a non-faulty system, missing data or disturbed or undisturbed measurements. In the general case, $\delta$ can take on $S$ different values. A common scenario is the binary case $S = 2$: disturbed/undisturbed, faulty/non-faulty etc.

To handle such systems, a Kalman filter bank can be used. It runs multiple filters in parallel that each have its own assumption of which mode $\delta_k$ the system is currently in.

A Kalman filter bank is a group of Kalman filters where each filter at each time instant has a parameter setup that is depending on a discrete mode, $\delta_k$. Consider the following jump Markov linear system

$$\begin{aligned} x_{k+1} &= A(\delta_k) + B(\delta_k)u_k + E(\delta_k)w_k \\ y_k &= C(\delta_k)x_k + D(\delta_k)u_k + e_k \end{aligned} \tag{2.41}$$

where $w_k$ is process noise distributed as $w_k \sim \mathcal{N}(0, Q_k(\delta_k))$ and $e_k$ is measurement noise distributed as $e_k \sim \mathcal{N}(0, R_k(\delta_k))$. The matrices $A(\cdot)$, $B(\cdot)$, $C(\cdot)$, $D(\cdot)$ and $E(\cdot)$ are known.

For a dynamic system, each time instant has a mode associated to it. That means

that for an experiment we will have an unknown mode sequence $\delta_{1:k}$ that is sought.

The estimated state sequence $x^i_{0:k}$ will be based on the assumed mode sequence $\hat{\delta}^i_{1:k}$. Each mode sequence has its individual combination hypothesis of what $\delta_{1:k}$ has actually been. Since the models are mode dependent, (2.41), different state sequences $x^i_{0:k}$ will be acquired for each mode sequence. The mode sequence $\delta_{1:k}$ to choose is the one that minimizes the prediction errors of the Kalman filter.

To explore the entire solution space, one should evaluate all possible mode combinations $\delta_{1:k}$. In the general case the number of combinations are $S^k$. This becomes unfeasible very soon since the computational complexity grows exponentially with time. Hence, approximations have to be made.

#### General Full Solution

The posterior distribution of the state can be expressed as a Gaussian mixture of all possible mode sequence realizations

$$p(x_k|y_{1:k}) = \sum_{i=1}^{S^k} p(\delta^i_{1:k}|y_{1:k}) p(x_k|\delta^i_{1:k}, y_{1:k}) \tag{2.42}$$

$$= \sum_{i=1}^{S^k} \mu^i_{1:k} \mathcal{N}(x_k; \hat{x}^i_{k|k}, P^i_{k|k}). \tag{2.43}$$

$\mu^i_{1:k} = p(\delta^i_{1:k}|y_{1:k})$ is subsequently a weight based on the probability of a certain mode sequence $\delta^i_{1:k}$ given the measurements $y_{1:k}$. For a specific mode sequence $i$, $\mathcal{N}(x_k; \hat{x}^i_{k|k}, P^i_{k|k})$ is straightforwardly calculated using a Kalman filter.

The minimum variance state estimate is therefore given as

$$\hat{x}_{k|k} = \sum_{i=1}^{S^k} \mu^i_{1:k} \hat{x}^i_k. \tag{2.44}$$

The weights $w^i_{1:k}$ are calculated for a certain sequence by recursively applying Bayes' theorem

$$\mu^i_{1:k} = p(\delta^i_{1:k}|y_{1:k}) = \frac{p(\delta^i_{1:k})}{p(y_{1:k})} p(y_{1:k}|\delta^i_{1:k}) \tag{2.45}$$

$$= \frac{p(\delta^i_k|\delta^i_{1:k-1}) p(\delta^i_{1:k-1})}{p(y_k|y_{1:k-1}) p(y_{1:k-1})} p(y_k|\delta^i_{1:k}, y_{1:k-1}) p(y_{1:k-1}|\delta^i_{1:k}) \tag{2.46}$$

$$= \frac{p(\delta^i_{1:k-1})}{p(y_{1:k-1})} p(y_{1:k-1}|\delta^i_{1:k-1}) p(\delta^i_k|\delta^i_{1:k-1}) \frac{p(y_k|\delta^i_{1:k}, y_{1:k-1})}{p(y_k|y_{1:k-1})} \tag{2.47}$$

$$= \mu^i_{1:k-1} \Pi_{(\delta^i_k, \delta^i_{k-1})} \frac{p(y_k|\delta^i_{1:k}, y_{1:k-1})}{p(y_k|y_{1:k-1})}. \tag{2.48}$$

$\Pi_{(\delta_k,\delta_{k-1})}$ is the mode switching probability that states what the probability is to switch from one mode to another. $\Pi_{(\delta_k,\delta_{k-1})}$ is assumed constant. The mode state is modeled as a time invariant Markov chain why the probability of mode switch is only dependent on the previous mode and not on the full mode history.

#### Implementation Approximations

A way of visualizing how the number of mode sequence combinations grows with time is to see it as a tree. For each new measurement, each combination $\delta^i_{1:k}$ is spawned off into $S$ more branches to provide all possible mode combinations.

Since the number of Gaussians in the mixture (2.42) grows exponentially as $S^k$, certain approximations must be applied to implement it. The two main approaches are pruning and merging.

Pruning means that mode sequences with a low probability are discarded. In the tree visualization, pruning means that low probably branches are simply cut off. If $p(\delta^i_{1:k}|y_{1:k})$ is very small, continually branching it is a waste of computational resources. For details on pruning and how to implement it, see Gustafsson (2010).

In merging, the branches are not cut off, they are joined with other branches, keeping the total number of branches constant. One can for example maintain the full mode sequence over a given time window $L$, $\delta^i_{k-L:k}$, and merge these branches when another measurement is available. A mixture of $N$ Gaussians can be approximated as

$$\begin{aligned} p(x_k) &= \sum_{i=1}^{N} \mu^i_{1:k} \mathcal{N}(x_k; \hat{x}^i_k, P^i_k) \\ &\approx \mathcal{N}(x_k; \hat{x}_k, P_k) \end{aligned} \tag{2.49}$$

where

$$\hat{x}_k = \sum_{i=1}^{N} \mu^i_k \hat{x}^i_k \tag{2.50}$$

$$P_k = \sum_{i=1}^{N} \mu^i_k \Big( P^i_k + (\hat{x}^i_k - \hat{x}_k)(\hat{x}^i_k - \hat{x}_k)^T \Big) \tag{2.51}$$

This approximation preserves the first and second moments.

### 2.3.5 Interacting Multiple Model

One of the most popular merging methods is the Interacting Multiple Model (IMM) filter Bar-Shalom et al. (2001). For a system with $S$ modes, it makes the approximation

$$p(x_k|y_{1:k}) \approx \sum_{i=1}^{S} \mu^i_k \mathcal{N}(x_k; \hat{x}^i_{k|k}, P^i_{k|k}). \tag{2.52}$$

Note that the posterior mode probabilities are $\mu_k^i = p(\delta_k^i|y_{1:k})$, why the entire mode sequence history has been discarded. In the IMM filter, only the probability of the current mode is maintained. The IMM filter is applied in Paper A.

### Filter Equations

In the IMM filter, the states $x_k^i$ with covariance $P_k^i$ are estimated, one pair for each mode, and the mode probabilities $\mu_k^i$, one for each mode, using the measurements $y_{1:k}$.

The algorithm of how to merge the estimates, the time update, the measurement update and the mode probability estimation are given below for convenience.

When a new time update is needed, the first step is to calculate the mixing probabilities $\{\mu_{k-1|k-1}^{ji}\}_{i,j=1}^{N_\delta}$

$$\mu_{k-1|k-1}^{ji} = \frac{\Pi_{ij}\mu_{k-1}^j}{\sum_{l=1}^{N_\delta}\Pi_{il}\mu_{k-1}^l}. \tag{2.53}$$

In doing so it uses the mode switching probabilities $\Pi_{ij}$ that states what the probability is of being in mode $i$, if one just previously was in mode $j$.

The mixed estimates $\{\hat{x}_{k-1|k-1}^{0i}\}_{i=1}^{N_\delta}$ and covariances $\{P_{k-1|k-1}^{0i}\}_{i=1}^{N_\delta}$ are then produced as

$$\hat{x}_{k-1|k-1}^{0i} = \sum_{j=1}^{N_\delta}\mu_{k-1|k-1}^{ji}\hat{x}_{k-1|k-1}^{j} \tag{2.54}$$

$$\begin{aligned} P_{k-1|k-1}^{0i} = \sum_{j=1}^{N_\delta}\mu_{k-1|k-1}^{ji}[P_{k-1|k-1}^{j} + \\ + (\hat{x}_{k-1|k-1}^{j} - \hat{x}_{k-1|k-1}^{0j})(\hat{x}_{k-1|k-1}^{j} - \hat{x}_{k-1|k-1}^{0j})^T] \end{aligned} \tag{2.55}$$

In this step, a bit of the other state estimates from the other modes are mixed in. How much is mixed in is based on the current mode probabilities and the mode switching probabilities. These mixed estimates are now time updated using

$$\begin{aligned} \hat{x}_{k|k-1}^{i} &= A_k^i\hat{x}_{k-1|k-1}^{0i} \\ P_{k|k-1}^{i} &= A_k^iP_{k-1|k-1}^{0i}A_k^{iT} + E_k^iQ_k^iE_k^{iT}. \end{aligned} \tag{2.56}$$

In both the time update and the measurement update, the model matrices associated to each corresponding mode is used. That means that when the state vector $x_k^i$ is updated, the state transition matrix $A_k^i = A_k(\delta^i)$ is used, for example.

The new measurement $y_k$ is used to update the states using the standard Kalman

filter measurement update.

$$\hat{x}^i_{k|k} = \hat{x}^i_{k|k-1} + K^i_k(y_k - C^i_k \hat{x}^i_{k|k-1}) \tag{2.57}$$

$$P^i_{k|k} = P^i_{k|k-1} - K^i_k S^i_k K^{iT}_k \tag{2.58}$$

$$S^i_k = C^i_k P^i_{k|k-1} C^{iT}_k + R^i_k \tag{2.59}$$

$$K^i_k = P^i_{k|k-1} C^{iT}_k (S^i_k)^{-1} \tag{2.60}$$

The mode probabilities $w^i_k$ are then updated using the previous mode probabilities, the mode switching probabilities and the innovation with innovation covariance for each filter.

$$\mu^i_k = \frac{\mathcal{N}(y_k; C^i_k \hat{x}^i_{k|k-1}, S^i_k) \sum_{j=1}^{N_\delta} \Pi_{ji} \mu^j_{k-1}}{\sum_{l=1}^{N_\delta} \mathcal{N}(y_k; C^i_k \hat{x}^l_{k|k-1}, S^l_k) \sum_{j=1}^{N_\delta} \Pi_{jl} \mu^j_{k-1}} \tag{2.61}$$

The updated mode probabilities $\mu^i_k$ are finally used to calculate the overall estimate $\hat{x}_{k|k}$ and covariance $P_{k|k}$ as

$$\hat{x}_{k|k} = \sum_{i=1}^{N_\delta} \mu^i_k \hat{x}^i_{k|k} \tag{2.62}$$

$$P_{k|k} = \sum_{i=1}^{N_\delta} \mu^i_k [P^i_{k|k} + (\hat{x}^i_{k|k} - \hat{x}_{k|k})(\hat{x}^i_{k|k} - \hat{x}_{k|k})^T]. \tag{2.63}$$

### 2.3.6 Hidden Markov Model

If a discrete mode $\delta_k$ is the only state that is sought, a Hidden Markov Model (HMM) can be used. In a system with $S$ possible modes, the HMM estimates the probability of being in each mode using a test statistic $\lambda_k$, and a mode switching probability $\Pi_{i,j}$. HMMs are used in Paper C.

The test statistic $\lambda_k$ can be constructed using the measurements $y_k$ and has a known distribution for each mode $\delta^i_k$. For example we could have $\lambda^i_k \sim \mathcal{N}(0, \sigma^2_i)$ if $\delta_k = i$ and $\lambda^j_k \sim \mathcal{N}(m_j, \sigma^2_j)$, if $\delta_k = j$. To make the modes easily distinguishable, we want $\sigma_j \neq \sigma_i$ and $m_j \neq 0$.

The mode switching probability $\Pi_{i,j}$ introduces a sort of dynamics to the system. Most often, $\lambda_k$ depends only on the latest measurement $y_k$, and $\Pi_{i,j}$ ensures that $\mu_k$ does not solely depend on $y_k$ but on previous measurements as well. $\Pi_{i,j}$ does this by stating the probability of switching from mode $j$ to mode $i$. So if $\lambda_k$ indicates that a very improbable mode switch has occurred, $\Pi_{i,j}$ states that the final mode probabilities should only indicate a minor change. In short, to facilitate an improbable mode switch, the test statistic likelihood for the new mode has to be very large or else multiple test statistics need to indicate that the switch has occurred for the mode probabilities to change completely.

When a new measurement $y_k$ is available, $\lambda_k$ is calculated as

$$\lambda_k = g(y_k). \tag{2.64}$$

The probability of being in mode $i$ at time $k$, $\mu_k^i = P(\delta_k^i|y_k)$, is calculated recursively

$$\begin{aligned}\mu_k^i &= P(\delta_k^i|y_k) \propto p(y_k|\delta_k^i)P(\delta_k^i|y_{k-1}) \\ &= p(y_k|\delta_k^i)\sum_{j=1}^{N_\delta}\Pi_{ji}\mu_{k-1}^j \qquad (2.65) \\ &= p(\lambda_k|\delta_k^i)\sum_{j=1}^{N_\delta}\Pi_{ji}\mu_{k-1}^j \qquad (2.66)\end{aligned}$$

since $y_k$ and $\lambda_k$ contain the same information. $\mu_k^i$ then has to be normalized. Hence we have

$$\mu_k^i = \frac{p(\lambda_k|\delta_k^i)\sum_{j=1}^{N_\delta}\Pi_{ji}\mu_{k-1}^j}{\sum_{l=1}^{N_\delta}p(\lambda_k|\delta_k^l)\sum_{j=1}^{N_\delta}\Pi_{jl}\mu_{k-1}^j}. \tag{2.67}$$

If the probability transition matrix is chosen as having the same probability of change between all modes, i.e.

$$\Pi^c = \begin{bmatrix} 1/S & \dots & 1/S \\ \vdots & \ddots & \vdots \\ 1/S & \dots & 1/S \end{bmatrix} \tag{2.68}$$

the probability of each mode is only depending on the latest measurement.

$$\begin{aligned}\mu_k^i &= \frac{p(\lambda_k|\delta_k^i)\sum_{j=1}^{N_\delta}\Pi_{ji}^c\mu_{k-1}^j}{\sum_{l=1}^{N_\delta}p(\lambda_k|\delta_k^l)\sum_{j=1}^{N_\delta}\Pi_{jl}^c\mu_{k-1}^j} \\ &= \frac{p(\lambda_k|\delta_k^i)\alpha\sum_{j=1}^{N_\delta}\mu_{k-1}^j}{\sum_{l=1}^{N_\delta}p(\lambda_k|\delta_k^l)\alpha\sum_{j=1}^{N_\delta}\mu_{k-1}^j} \\ &= \frac{p(\lambda_k|\delta_k^i)}{\sum_{l=1}^{N_\delta}p(\lambda_k|\delta_k^l)} \qquad (2.69)\end{aligned}$$

All dynamics are then removed and only the latest measurement determines the mode. This makes it prone to give an erroneous mode classification after receiving just one troublesome measurement.

## 2.4 Estimation under Disturbances

This section discusses a fundamental estimation problem that occurs in many shapes, and has certain applications in this thesis. The situation occurs when

there are two related measurements. The first sensor measures a physical quantity with a large deterministic disturbances, which can hardly be modeled as a stochastic process. The second sensor measures the derivative of the physical quantity, but with a small offset. Thus, integrating the second sensor corresponds to the physical quantity but with a linear drift over time.

### 2.4.1 Problem Fundamentals

We want to estimate a feature $p_k$ of a dynamic system like a position or a heading. The feature is measured, $y_k^p$, but the measurements suffer from deterministic disturbances $d_k^p$, like location dependent disturbances.

Available is also a measurement of the system movements, like the derivative of $p_k$, called $y_k^v$. This measurement is unaffected by $d_k^p$ but the signal is corrupted by noise and slight sensor deficiencies like for example a slight bias. These errors are of the types that cannot be determined once and for all in a calibration procedure, since they might change slowly over time or be different every time the sensor is started.

We assign a state $v_k$ for the derivative state to get the dynamic model

$$\begin{pmatrix} p_{k+1} \\ v_{k+1} \end{pmatrix} = \begin{pmatrix} 1 & T \\ 0 & 1 \end{pmatrix} \begin{pmatrix} p_k \\ v_k \end{pmatrix} + \begin{pmatrix} T^2/2 \\ T \end{pmatrix} w_k \tag{2.70}$$

where $w_k^p \sim \mathcal{N}(0, Q_k)$. We also have the the measurement models

$$y_k^p = h(p_k) + d_k^p + e_k^p \tag{2.71a}$$
$$y_k^v = g(v_k) + e_k^v \tag{2.71b}$$

where $e_k^p \sim \mathcal{N}(0, R_p)$ and $e_k^v \sim \mathcal{N}(0, R_v)$.

Now, if $d_k^p = 0$ a nonlinear filter can be applied to estimate $p_k$ and $v_k$. If $h(\cdot)$ and $g(\cdot)$ are linear a Kalman filter can be used.

If all measurements $y_{1:k}^p$ are corrupted by large deterministic disturbances, i.e. $d_{1:k}^p \neq 0$, and $p_0$ is unknown, $p_k$ cannot be estimated correctly. This is natural since if all measurements of position is bad, estimating position using only velocity is impossible.

The question is how to estimate $p_k$ and $v_k$ when some $d_i^p \neq 0$?

The estimation system should first and foremost be robust against the disturbances $d_k^p$. By robust we mean that $\hat{p}_k$ should not be pulled away by the disturbances. Also, during the disturbed periods, the filter has to rely on the dynamic model only, but despite errors in $y_k^v$, $\hat{p}_k$ should not drift away. In short, we want the filter to be both resistant to disturbances and movement measurement deficiencies.

### 2.4.2 General Solution Outline

When the disturbances come and go, $p_k$ is observable when $d_k^p = 0$ if we also know that $d_k^p = 0$. So in order to utilize $y_k^p$ in our estimation system at all, we need to detect when $d_k^p = 0$ and when $d_k^p \neq 0$. Given that we know or have an estimate of when $d_k^p = 0$, $p_k$ can be estimated for those time instances.

During the disturbed periods, we have to rely on the dynamic model (2.70) and (2.71b). Unfortunately, errors in $y_k^v$ will cause $\hat{p}_k$ to drift off. To be able to handle long periods of disturbances, the sensor inaccuracies corrupting $y_k^v$ will therefore also have to be modeled and estimated.

We assume that the sensor deficiencies can be modeled using the constant error parameters $\theta$. We extend the state vector to also contain the assumed constant fault parameter vector

$$\begin{pmatrix} p_{k+1} \\ v_{k+1} \\ \theta \end{pmatrix} = \begin{pmatrix} 1 & T & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} p_k \\ v_k \\ \theta \end{pmatrix} + \begin{pmatrix} T^2/2 \\ T \\ 0 \end{pmatrix} w_k \tag{2.72}$$

$$y_k^p = h(p_k) + d_k^p + e_k^p \tag{2.73a}$$
$$y_k^v = g(v_k, \theta) + e_k^v \tag{2.73b}$$

If the error modeling is feasible and $\hat{\theta}$ is estimated correctly, this model is a better description of the true system. It can therefore be relied upon during much longer periods of disturbance. In most practical applications though, eventually $\hat{p}_k$ will start drifting off.

So, not only do we have to estimate $p_k$ and $v_k$, we also have to estimate $\theta$. Given that we know when the measurements $y_k^p$ are disturbed and when they are not, the undisturbed sections of $y_k^p$ can be used to estimate $\theta$ and $p_k$.

Since we do not know if $d_k^p = 0$ or $d_k^p \neq 0$, this has to be estimated as well. In the end, we have to estimate $p_k$, $v_k$, $\theta$ and the system mode, i.e. if $d_k^p = 0$ or $d_k^p \neq 0$.

### 2.4.3 State and Mode Estimation

This type of problem fits into the mode based estimation systems such as a Kalman filter bank, Section 2.3.4, or an HMM filter, Section 2.3.6.

In a Kalman filter bank solution, the system switches between different models depending on the estimated mode state. The mode estimate depends on the likelihood of the measurement given the mode and the previous one, $p(y_k|\delta_{1:k}^i, y_{1:k-1})$. In the IMM filter, this means that the mode estimation is based on the innovation of each filter, i.e. $\tilde{y}_k^{p,i} = y_k^p - h(\hat{p}_k^i)$.

The advantage of estimating state and mode jointly using an IMM filter, is that

when the filter has converged, detecting if $d_k^p \neq 0$ is very straightforward. Unfortunately, if the filter has diverged, all measurements $y_{k:N}^p$ will be erroneously interpreted as disturbed. The filter has to be reinitialized to function properly again.

Using an HMM, filter mode probabilities $\mu_k^p$ can be estimated using a test statistic independent of the states. A test statistic can be constructed as

$$\lambda_k = f(y_k^p, y_k^v) \tag{2.74}$$

It is quite common that the distribution of the test statistic is known in the usdisturbed case but unknown for the disturbed case, which has to be approximated.

Once the state mode has been estimated, the states can be estimated using a nonlinear filter where for example $R_k^p = R^p(\mu_k)$.

If the vector of estimates $\hat{p}_{0:k}$ is sought and the system has linear dynamics, the state estimation can often be formulated on batch form, see Section 2.3.3. Again, the measurement noise $R_k^p$ can be made mode dependent. Using a batch formulation, the smoothing estimate of $\hat{p}_{0:k-1}$ and the filtering estimate of $\hat{p}_k$ is acquired.

The number of states can be reduced by using a dynamic model on input form where $y_k^v$ drives the time update.

$$p_{k+1} = p_k + f(y_k^v, \theta) + e_k \tag{2.75}$$

Such a model can be approximated as a function of only $p_0$ and $\theta$, see Section 2.2.3. In this thesis this has been applied on the heading estimation problem, see Papers B and C.

### 2.4.4 Examples of Problems

The problem described above is quite common in practical sensor fusion applications. The simplest case is the scalar one where $p_k$ is one-dimensional.

- Position estimation on rails or roads from a speedometer, which is reliable and accurate up to scale and bias, and position measurements from for example GPS or landmark tracking. The position estimates are uncertain due to satellite occlusion or lack of steady landmarks.
- Speed estimation for an autonomous underwater vehicle (AUV). $p_k$ is velocity, accelerometer measurements $y_k^v$ are used as input signal and $y_k^p$ are measurements from an unreliable speedometer, e.g. Doppler velocity.
- Heading estimation where $y_k^p$ are either heading measurements from a magnetometer or difference-of-arrival (DOA) from radio transmitters estimated by an antenna array. Both measurement types are prone to large and frequent disturbances. The DOA estimates suffer from multipathing, where the angle may correspond to the position of strong reflectors in the environment, see Figure 3.2 from Erkstam and Tjernqvist (2012) for an example. Magnetometers are often disturbed by location dependent interferences indoors, see Paper A. To support the heading estimation, a gyro can be used

as $y_k^v$ but such measurements often have bias and scale errors. The heading estimation problem using a disturbed magnetometer and a low grade gyro will be studied in depth in this thesis, see Section 3 and Papers A, B and C.

A two dimensional example of such a system is positioning of a car with GPS and wheel encoders. The GPS measurements are often good but can be disturbed by large buildings or tunnels. The wheel encoders measure how the wheels are turning, thereby giving estimates of the velocity and turning rate of the vehicle. Unfortunately these movement estimates will have errors due to wheel slip, sensor errors, slightly different wheel diameters etc. So in this scenario, $p_k$ is vehicle position and heading, $y_k^p$ is the GPS measurement and $y_k^v$ are the wheel encoders. If the road network is known, the problem complexity can be reduced significantly.

### 2.4.5 Discussion

On most occasions the measurement function $h(\cdot)$ is just a measurement of the state, $h(p_k) = p_k$. In the case of heading estimation using a magnetometer it is instead $h(p_k) = p_k \mod 2\pi$. How to handle this in the batch estimation case is discussed in Paper B.

The question of using the full state model (2.72) or the input model (2.75) is to a large extent a question of how many states one prefer. As was shown in Callmer (2011), the full model introduces a time delay in the system dynamics and also a low pass filtering of the noise in $y_k^v$. If $y_k^v$ does not suffer from outliers, missing data or similar, the input form is the more attractive one because of the delay and lower state dimension.

If the dynamic model (2.75) is also on the reduced form presented in Section 2.2.3, the only unknown states are $p_0$ and $\theta$. Estimating only these few constant states using $y_{1:k}^p$ and $y_{1:k}^v$ on batch form may seem like quite an overkill, but the resulting estimates $\hat{p}_{0:k}$ will be very resistant to both disturbances and sensor errors. More on this in Papers B and C.

# 3

# Indoor Positioning

The ability to position people or equipment indoors has become a very hot topic in the last couple of years. What used to be a mostly academic or military field of interest has now also entered the commercial sector. Along the way one thing has been lost though: a clear understanding of what we are actually talking about.

The question has become if one can solve the indoor positioning problem (see for example Economist (2012)) but it is not a relevant question. The question in general has no answer since the solution relies heavily on a different question: for whom or what it shall be solved and for what purpose? There is no silver bullet that solves the entire problem, but there are many different types of solutions, each working in its own scenario.

The broadest way to divide the indoor positioning issue into smaller manageable ones, is to start with if the indoor environment is known or unknown. There is a huge difference between positioning a soldier inside an unknown building in an unknown country, or to position ones own things in ones own building.

If one has control over the environment, two things in particular can be utilized to aid in the positioning. First of all maps of the building can be utilized. They can restrict the possible user movements, thereby significantly lowering the uncertainty in the position estimate. The second thing one can do is to install equipment that aids the positioning such as radio beacons, tags for cameras or reflectors for scanners. Using these, one can produce environmental maps that show what will be measured where, which will provide great support for a positioning system.

The other indoor positioning problem is the one in which we have no control over the environment. The positioning should be performed in an unknown building

where the only information available is from sensors carried by the subject. This makes this problem scenario much more challenging.

Again, how to solve this problem depend on the reason why one wants it to be solved. If the device used is a robot, one can for example equip it with wheel encoders for odometry and cameras and laser scanners to support the positioning. In such a scenario the sensor platform is fixed making the geometrical relation between the sensors known. If the end user is a human such as a first responder or a soldier, there are limits to which sensors can be carried without them affecting the user's ability to perform his or her task. Also, to keep the cost of the system low, the choice of sensor often becomes a cost/performance tradeoff.

The problem of indoor positioning for humans in an unknown environment is the primary one studied in this thesis. Such systems are often based on a dead reckoning system estimating the step-to-step movements. In this chapter we will give a background to that problem and then describe a part of the dead reckoning system in a bit more detail. We will also give a short presentation of additional techniques in the literature that have been used to support the dead reckoning system, utilizing for example information about the building.

## 3.1 Human Positioning in an Unknown Environment

In some scenarios, the operational environment is more or less completely unknown. The chances of there being a digital map available of a building in a warzone is slim at best. Maps of residential homes are also unlikely to be available to aid positioning for police or fire services, for the next couple of years. A positioning system for such environments will therefore have to be based on other techniques.

A basic system often used in indoor positioning is the pedestrian dead reckoning system (PDR). It estimates user position relative to a known starting position using the estimated user movements. User movements are measured using an IMMU and the accelerations measured are translated into a change in position using double integration. For such a system to give useful position estimates, the sensor is often mounted on the foot. By detecting every time the foot is at rest, the accumulated estimation errors can be corrected. For descriptions of pedestrian dead reckoning systems, see for example Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007); Godha et al. (2006); Woodman and Harle (2009); Jiménez et al. (2010a); Bebek et al. (2010). Such a system will also be briefly described in Section 3.3 and a stand still detection system is presented in Section 3.3.1.

Using such a system, the distance traveled can be estimated well but there is a drift in heading causing positioning errors. In this thesis the problem of heading estimation is given a lot of attention, Papers A and B, and when such a system is included in a positioning system, the performance increases dramatically. This is shown in Paper C. Our system is still just a dead reckoning system. To achieve

even better positioning, more information is needed.

A well studied problem in autonomous robotics is the simultaneous localization and mapping (SLAM) problem. It is the problem of creating a map of the surroundings while at the same time positioning oneself in the very same map. A tutorial on the subject is Durrant-Whyte and Bailey (2006); Bailey and Durrant-Whyte (2006). In SLAM, a key to reducing the errors in estimated position is the so called loop closure. If one can detect that an area is being revisited, the error and uncertainty in estimated user position can be reduced significantly.

To produce loop closures in indoor positioning for humans, for example a camera can be used, Rydell and Andersson (2010). If one does not want to add more sensors, user activities can instead be used as loop closure features. In Grzonka et al. (2010), the 'opening door' activity was detected and such events were connected to each other. By estimating also the number of doors and their specific locations, 'door opening' events could be used as loop closures. This was shown to significantly reduce the errors in estimated user position. Such passive events could be replaced by active events like user induced markers in the data. For example, if one believes the area one is in will be revisited, a marker in the data, like a special sensor wiggle, could be added. By connecting those markers, the same effect can be achieved.

Loop closures can also be made by having the user place out radio beacons along the path, Renaudin et al. (2007). The next time the beacon is passed, a loop closure is called and the built up errors are reduced. More on radio based positioning below.

In Angermann and Robertson (2012), SLAM was performed using only the user movements. By assuming that the user walk the same paths multiple times, the paths can be connected. When backtracking a path, assuming that the two paths are one and the same reduces the drift. That way biases in the system can also be estimated.

Even if one does not have a map of the building, some information about it is often still available. In Abdulrahim et al. (2011); Borenstein and Ojeda (2010) the user is assumed to primarily only walk in the principal directions of the building. For any given building, and especially office buildings, the main directions one can walk indoors are aligned with the outer walls. These four directions, sometimes expanded to eight, are the directions one can assume the user will most often be walking in. The estimated heading is therefore pushed towards the nearest principal direction.

Only knowing the outer walls of a building can also reduce the accumulated positioning errors, Beauregard et al. (2008). How much information such a map will provide is though strongly correlated to the layout of the building. In that work, the building was L-shaped and narrow, providing quite a lot of information. If the building is large and square, the positioning improvements will most likely be marginal.

## 3.2 Human Positioning in a Known Environment

When one is in control of the building or at least has a solid knowledge of it, positioning can be aided by utilizing this information. A map can be used to put restrictions on where it is actually possible to walk, which can provide a lot of support to a dead reckoning system.

To get actual measurements of absolute user positions, location dependent features like visual tags, Zachariah and Jansson (2012), or radio environment measurements can be used. Knowing which wifi access points that can be seen where, provides a lot of support to the positioning system.

For many more techniques available for various kinds of indoor positioning, see Muthukrishnan et al. (2005); Gu et al. (2009).

### 3.2.1 Map Matching

The most common approach in map matching is to let the estimate of user position roam free as long as it does not try to walk through a wall.

In Woodman and Harle (2008), a PDR system is fused with map matching using a particle filter. The map is represented as a 2.5D set of polygons that have either passable or un-passable edges. Also stairs are included in the same manner, linking the floors. This representation is shown to significantly reduce the error in position, providing a drift free estimate in user position once the filter has converged.

Widyawan et al. (2007) also applies a PDR system to a map matching system using a particle filter. The map representation is here based directly on the walls, giving zero weight to particles that want to walk through a wall. A similar approach was taken in Krach and Robertson (2008).

A problem with such a system is that a position estimate in a large room will be given a higher weight, and therefore be perceived as more likely, than a position estimate in a more closed space, regardless of actual user position. Since the only thing limiting the system is walls, position estimates where the particles have a high risk of running into the walls, will be bad ones while the open areas will be good ones. One can therefore experience depletion where particles in closed spaces are resampled into open spaces even if the closed space area is the one the user is actually in.

In Kaisera et al. (2013) this was handled by weighting the particles using an angular probability density function that specifies some directions as more probable than others. Particles travelling along a corridor then get a higher weight than ones that will hit the wall soon.

A map aided dead reckoning system restricts the possible movements of the user but does not provide any measurements of absolute position. If the position estimates have diverged, only the uniqueness of the path the user has travelled can make it possible to rediscover the true position.

### 3.2.2 Radio Positioning

To get a measurement of the absolute position of the user, other features like the radio environment can be used. For example the wifi coverage in a building can be used to greatly restrict the number of possible user positions if measured.

Radio based positioning is an entire research area in itself and will here only be covered briefly on a more conceptual level. For more details, Bahl and Padmanabhan (2000); Li et al. (2006); Roos et al. (2002); Youssef and Agrawala (2005); Honkavirta et al. (2009) are a good start.

There are two main methods of getting a position measurement from the radio environment: model-based and map-based. Model-based approaches estimates the distance to access points in known locations using radio propagation models and the measured drop in signal strength. Map based methods use a radio map that for a given position specifies which signals that should be measured and what strengths these signals should have. The actual radio signals measured by the user can then be compared to the map to get a measurement of user position. The more radio signals and the more the indoor environment makes them change, the more accurate the position estimate will be. Unfortunately, it can be hard to get a good positioning performance, since also such a simple thing as the user orientation will affect the result, Dil and Havinga (2010)

Model-based approaches are cheap since the site need not be visited. Known access point locations and a map can be used to produce an approximate model. The downside is that the model is just approximate if the signals have to travel through walls. The location of the access points must also be known and so also the broadcasted signal strengths.

For the map-based approaches it is the complete opposite. The environment is often mapped manually, ignoring the problems of radio propagation models and knowing the access point locations. On the other hand, the map creation part is often tedious and expensive and if the access points move, the mapping has to be remade.

In Woodman and Harle (2009), wifi fingerprinting was used to initialize the user position and support the PDR/map matching positioning system. Similar approaches were taken in Seitz et al. (2010); Wang et al. (2007) with equally successful results. As long as the radio map is accurate, a highly robust position estimate is achieved.

If there are multiple users and few beacons, collaborative localization can improve the positioning significantly by measuring and sharing the distances between the users, Wymeersch and J. Lien (2009); Rantakokko et al. (2011).

## 3.3 Foot Mounted IMMU for Dead Reckoning

The pedestrian dead reckoning system mentioned above has become a very popular system and will most likely be the foundation of future commercial indoor

navigation systems. It is cheap, hard to jam, easy to operate and gives quite good positioning results.

The IMMU mounted on the boot has three dimensional accelerometers, gyroscopes and magnetometers to measure the movements of the users. Double integrating the accelerometer measurements gives user position relative to the starting point if the gravity component can be removed correctly. Since sensor orientation is estimated using semi grade gyros, the orientation estimate will be a little bit off. This results in parts of the gravity component to be interpreted as user accelerations, resulting in an ever growing position error.

By utilizing the stand still phases, i.e. detecting and exploiting when the foot is on the ground, accumulated velocity errors can be corrected. This so called zero velocity update significantly improves the position estimates, and reduces the error in distance travelled from cubic in time to linear.

To have robust zero velocity updates, we need robust zero velocity detection. We have therefore derived a stand still detection system with very few false positives.

The remainder of this chapter is a presentation of our probabilistic stand still detection system. The system is described in detail and its performance is evaluated for a couple of different sensor positions. This dead reckoning system is described in Paper C, so to avoid too much overlap it is not presented here.

### 3.3.1 Stand Still Detection

A reliable stand still detection system is key to creating a good dead reckoning system. Since most detection systems in the end become a tradeoff between missed detections and false positives, for a system like this, avoiding false positives is much more important than a few missed detections. The main purpose of the stand still updates is to anchor the estimates at a regular basis, clearing the built up errors in estimated velocity. Missing one or two stand stills here and there is therefore not a big problem.

A stand still detection framework can be based on two types of sensors. Either one uses special sensors just to detect that the foot is touching the ground, or one uses the IMMU to detect that the sensor is at rest. Theoretically, the former approach could provide safer stand still detections, but there are also some drawbacks. Often sensors designed to detect contact have some moving parts which are prone to be worn out by usage. Also, if the user is crawling, the foot is sometimes at a stand still even though the sole is not touching the ground. That means, to detect different types of standstill, multiple contact detecting sensors are needed on the boot. And finally, the fewer the sensors, the less that can go wrong. If one can detect the stand stills using the IMMU data itself, it is therefore better.

IMMU based stand still detection has most often been done by comparing the accelerometer or gyro signal to a threshold. Here we put the stand still detection in a probabilistic framework using a test statistic with known distributions and an HMM. The result is an estimated probability of stand still that is consecutively

calculated at every time instant. This can be used for zero velocity updates in a filtering framework.

Parts of our probabilistic framework for stand still detection have previously been presented in Callmer et al. (2010b), Callmer (2011) and Rantakokko et al. (2011). A very compressed version is also given in Paper C.

### Related Work

The most common approach to stand still detection is to use averaged accelerometer or gyro measurement and compare it to a threshold; Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007). The threshold is chosen *ad hoc* and is normally quite restrictive to minimize false positives.

Another approach is the moving variance used in Godha et al. (2006) where the variance computed over a sliding window is compared to a threshold. One of the problems with this approach is that in order to make it work properly, the time interval the boot needs to be still is quite long. The system can therefore work quite well if the user is walking but not so much during running.

Probabilistic zero velocity detection has previously been proposed in Skog (2009) and Skog et al. (2010) where a hypothesis test was used to determine if the foot was stationary or moving. The hypothesis test was performed using a test statistic based on a Generalized Likelihood Ratio Test (GLRT). The test statistic had a known stand still distribution and an unknown movement distribution that was approximated as a constant. Since one of the probability density functions was constant, the GLRT test boiled down to comparing a test statistic of the norm of the acceleration measurements and the angular velocity measurements, to a threshold. If the test statistic was smaller than the threshold, a stand still was declared. Since the test statistic had an unknown movement distribution, the threshold was chosen *ad hoc*, making the framework in practice similar to the ones in Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007).

### System Outline

The stand still detection system is mode based and uses an HMM filter to estimate the probability of stand still and movement, respectively, Section 2.3.6. As input, the HMM uses a test statistic based on the IMMU signals.

The test statistic is based on the latest gyro and accelerometer measurements and has a known distribution during stand still, since the sensor should then be at rest. The IMMU was mounted by the shoe laces to get a good tradeoff between stand still detection performance and gyro dynamic range.

**Test Statistics Derivation** An example of a walking sequence with a shoe lace mounted IMMU is shown in Figure 3.1. The foot is stationary around sample 530, 650, 770, 880 and 980. During these phases the norm of the accelerometer signals is the gravitation constant with some added noise. At the same time the norm of the angular velocity signal is zero with some additive noise. These are the key

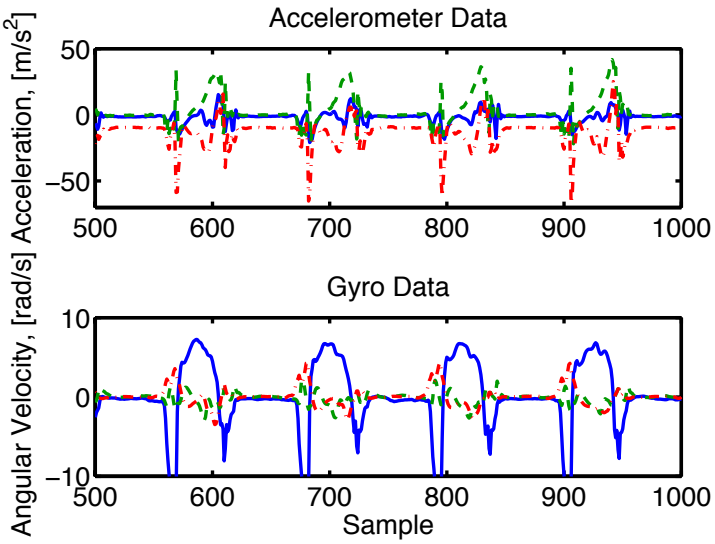


***Figure 3.1:*** *Example of accelerometer data (where $x$, $y$ and $z$ is solid blue, dashed green and dashdotted red, respectively) and gyro data ($\omega_x$, $\omega_y$ and $\omega_z$ is solid blue, dashed green and dashdotted red, $\omega_i$ is angular rotation rate around axis $i$) during a walking sequence. The foot is stationary around time instances 530, 650, 770, 880 and 980.*

characteristics of the test statistic.

**Sensor Models** The signal model is

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{y}_k^a(\delta_k) \\ \mathbf{y}_k^\omega(\delta_k) \end{bmatrix} + \begin{bmatrix} \mathbf{v}_k^a \\ \mathbf{v}_k^\omega \end{bmatrix} \tag{3.1}$$

where $\mathbf{y}_k^a$ and $\mathbf{y}_k^\omega$ denote the acceleration measurement vector and the angular velocity measurement vector, respectively. Further, $\delta_k$ denotes the model dependence on the phase of the human step sequence at time $k$. Naturally, the model differs significantly between when the foot is at stand still and when it is moving.

The measurements are assumed to have additive independent identically distributed Gaussian noise $\mathbf{v}^a \sim \mathcal{N}(\mathbf{0}, \sigma_a^2)$ and $\mathbf{v}^\omega \sim \mathcal{N}(\mathbf{0}, \sigma_\omega^2)$ where $\sigma_\omega^2 = \sigma_\omega^2 \mathbf{I}$ and $\sigma_a^2 = \sigma_a^2 \mathbf{I}$ and $\mathbf{I}$ is the $3 \times 3$ identity matrix.

During stand still the sensor model is

$$\begin{bmatrix} \mathbf{y}_k^a \\ \mathbf{y}_k^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u}_k \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{v}_k^a \\ \mathbf{v}_k^\omega \end{bmatrix}, \tag{3.2}$$

where $\mathbf{u}_k$ is the unknown gravitational direction vector and $g$ is the gravitational constant 9.81. Since the orientation of the boot changes over time, so does $\mathbf{u}_k$.

When the foot is moving the sensor model changes to

$$\begin{bmatrix} \mathbf{y}_k^a \\ \mathbf{y}_k^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u}_k + \mathbf{a}_k \\ \boldsymbol{\omega}_k \end{bmatrix} + \begin{bmatrix} \mathbf{v}_k^a \\ \mathbf{v}_k^\omega \end{bmatrix} \tag{3.3}$$

where $\mathbf{a}_k$ is user acceleration and $\boldsymbol{\omega}_k$ is user angular velocity which therefore have unknown distributions.

**Test Statistics** We construct a test statistic by combining the acceleration and angular velocity measurements

$$\lambda_k = \frac{\|\mathbf{y}_k^a\|^2}{\sigma_a^2} + \frac{\|\mathbf{y}_k^\omega\|^2}{\sigma_\omega^2} \tag{3.4}$$

where $\lambda_k \sim \chi^2(6, \lambda)$ during stand still. It has a non-central chi-square distribution since $\mathbf{y}_k^a$ has nonzero mean when the foot is stationary. Its non-centrality parameter $\lambda = g^2/\sigma_a^2$ and 6 is the number of degrees of freedom.

**Test Statistic Appearance during Walking Sequence** The test statistic for the walking sequence in Figure 3.1 is plotted in log-scale in Figure 3.2. The mean of the stand still distribution is marked with a dashed line. The stand still events occurring around time instances 530, 650, 770, 880 and 980 are clearly visible.

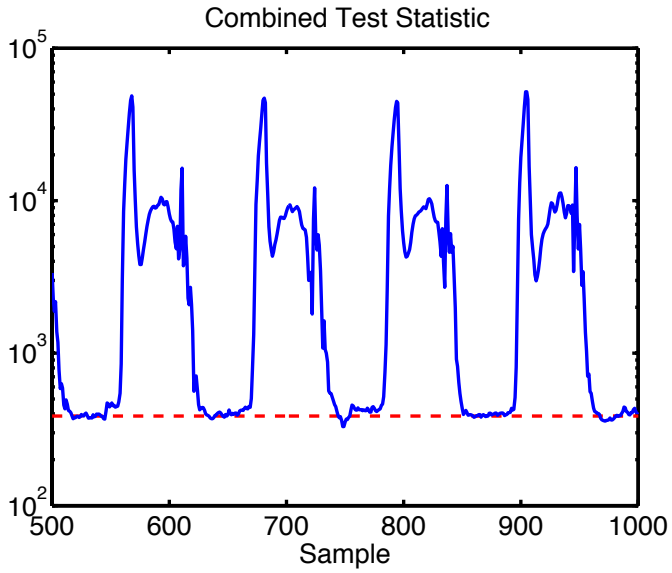


***Figure 3.2:*** *Logarithmic plot of the test statistic with the mean of the stand still distribution marked with a dashed line. The foot is stationary around time instances 530, 650, 770, 880 and 980.*

$\lambda_k$ has a movement distribution that is significantly larger than the stand still distribution. The stand still phases are all close to the mean of the stand still distri-

bution. Considering the significant spread between the stand still and movement distributions, detecting the stand still phases should be quite straightforward.

#### Test Statistic Distribution Validation

The test statistic stand still and movement distributions must be validated using experimental data. We need to make an approximation of the movement distribution for the HMM and also to ensure that the experimental stand still distribution is similar to the theoretical ones. The latter are plotted in Figure 3.3 where the empirical stand still distribution has been approximated using a histogram. There is a slight offset between the theoretical distribution and the empirical distribution. This is mainly related to how the data for the empirical evaluation was chosen. If the stand still data samples are chosen more conservatively, the mean of the empirical distribution will decrease.

The empirical movement distribution is shown as a histogram and is plotted with its approximation, Figure 3.4. The movement distribution was approximated using two Gaussians as

$$p_m(\lambda) = 0.92 \cdot N(3000, 7000) + 0.08 \cdot N(600, 150).$$

The histograms were created using a large amount of experimental data. Figure 3.5 show the histograms of both the stand still data and the experimental data. Clearly the two distributions are well separated.

The same movement distribution can be used for other types of movement as well like running, since the differences between different movement distributions are quite small, especially in the border region close to the stand still distributions.

#### Mode Probability Estimation

To determine the probability of stand still, an HMM is used, Section 2.3.6. In this case it has two modes: stand still and movement, and it estimates the probability of being in each mode at every time instant.

The mode transition probability matrix states the probability of switching modes which introduces some dynamics into the mode estimation. A lower mode transition probability requires a measurement with a higher likelihood for a mode switch to occur and vice versa.

The mode transition probability matrix used in the experiment is

$$\Pi = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix} \tag{3.5}$$

which states that the probability of going from stand still to moving or vice versa, is 2%. During normal walking the right foot takes about one step per second which results in roughly 2 mode transitions every 100 measurements.

The mode probabilities at time $k$ are calculated using the recursion (2.67). The probability density function of movement used in the HMM, is an approximation set to resemble the empirical movement density function, Figures 3.4.

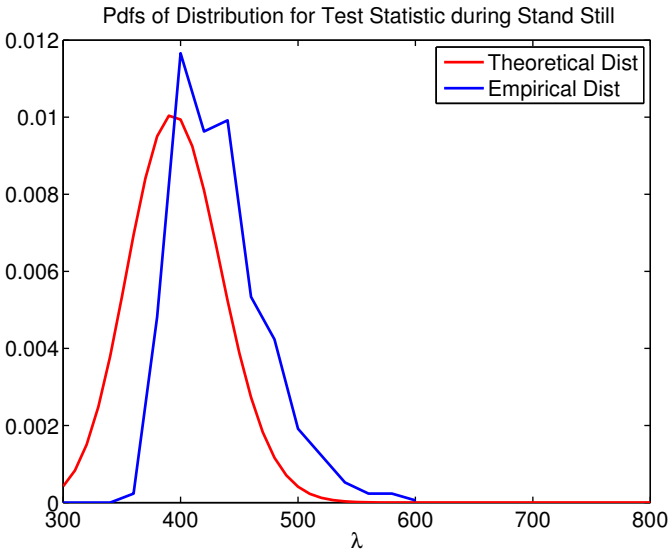


***Figure 3.3:*** *Theoretical stand still distribution of $\lambda$ and histogram of stand still measurements of $\lambda$ from experimental data. Offset is mainly due to how empirical data was chosen.*

For the binary mode case stand still and movement, the algorithm can be summarized as in Algorithm 4.

### Experimental Results

The mode estimation system was evaluated on the data sequence in Figure 3.1. The resulting mode probabilities provided by the HMM are shown in Figure 3.6. The measurement noises were set as $\sigma_a = 0.5$ and $\sigma_\omega = 0.08$, estimated from data when the sensor was at rest.

The stand still detection framework has no trouble detecting the stand still phases, Figure 3.6. The movement probability is very close to 1 during the movement phases while the stand still probability is about 0.999 when the foot is at rest. The stand still probabilities fluctuate a bit because they are quite sensitive to movements. Since a foot mounted IMMU is not entirely still just because the foot is touching the ground, this is entirely natural.

Longer walking experiments with 318 steps reveal that all 318 stationary phases were detected with no false positives. In these experiments a winter boot with a toe mounted sensor was used.

One of the big problems when evaluating indoor navigation systems is how to acquire reliable ground truth data. You can count the number of steps you take but that only works when you are just walking straight ahead and that is the least interesting case since it is the simplest. A data set consisting of hundreds of steps

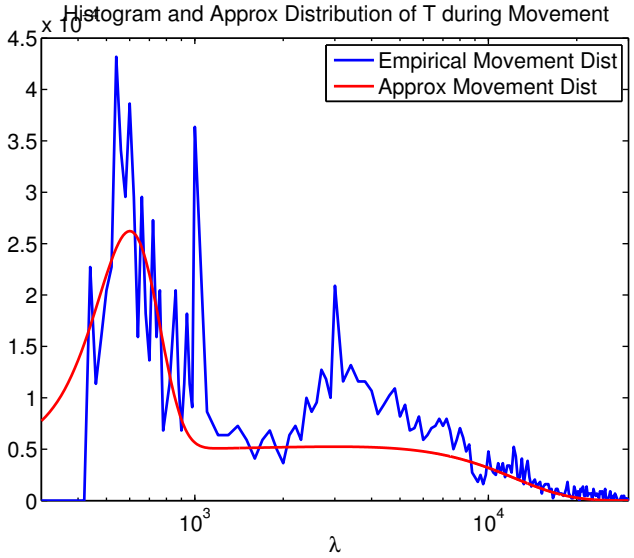


***Figure 3.4:*** *Histogram of movement measurements of $\lambda$ from experimental data, and the approximated Gaussian sum that resemble it.*

when you open and close doors, move around furniture and such, is much more difficult to analyze afterwards to detect exactly when the foot was still. One way to evaluate stand still detection performance is to incorporate it in a positioning system. If the positioning works well, the stand still detection is at least not working poorly.

### Conclusions

A test statistics with known stand still distribution has been evaluated for stand still detection. In conjunction with an HMM, the mode probabilities are readily calculated and can be used for zero velocity updates. The framework has been shown to safely detect stand still phases in the data.

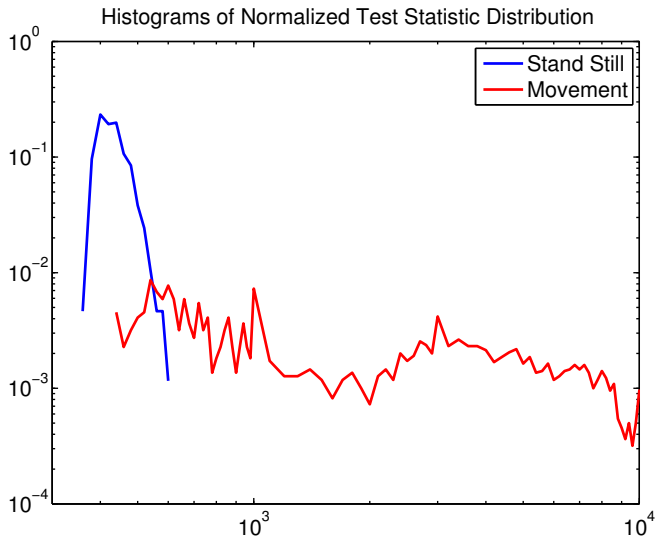


***Figure 3.5:*** *Histograms of stand still and movement data of $\lambda$ from experimental data. The two different types of movements are well separated.*

**Algorithm 3** Stand Still Detection Framework

**Require:** Stand still distribution $p_s$ and movement distribution $p_m$ of test statistic $T$. Measurements $\mathbf{y}^a$ and $\mathbf{y}^\omega$ with noise parameters $\sigma_a$ and $\sigma_\omega$. Mode transition probability matrix $\Pi$.

1: **for** $k = 1, \ldots, N$ **do**

2: Compute test statistic.

$$\lambda_k = \frac{\|\mathbf{y}_k^a\|^2}{\sigma_a^2} + \frac{\|\mathbf{y}_k^\omega\|^2}{\sigma_\omega^2} \tag{3.6}$$

3: Estimate stand still and movement probabilities, $\mu_k^s$ and $\mu_k^m$.

$$w_k^s = \frac{P_{s,k}}{P_{s,k} + P_{m,k}} \qquad w_k^m = \frac{P_{m,k}}{P_{s,k} + P_{m,k}}$$

where

$$\begin{aligned} P_{s,k} &= p_s(\lambda_k)\left(\Pi_{ss}\mu_{k-1}^s + \Pi_{sm}\mu_{k-1}^m\right) \\ P_{m,k} &= p_m(\lambda_k)\left(\Pi_{ms}\mu_{k-1}^s + \Pi_{mm}\mu_{k-1}^m\right) \end{aligned} \tag{3.7}$$

4: **end for**

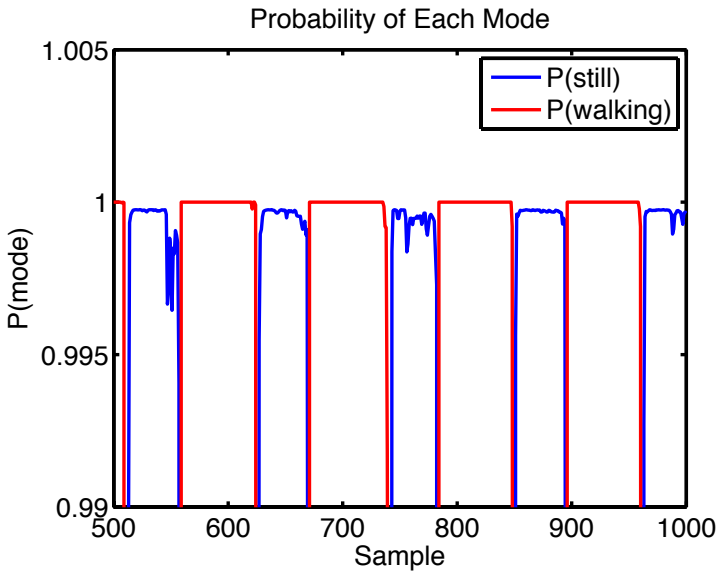


***Figure 3.6:*** *Mode probabilities from the data set in Figure 3.1. The foot is stationary around samples 530, 650, 770, 880 and 980.*

### 3.3.2 Stand Still Detection Performance for Different IMMU Positions

For a localization framework with a foot mounted IMMU, the sensor must be mounted on the shoe. Where the IMMU is attached on the boot affects the performance, not only of the stand still detection but also of the positioning. For stand still detection, the sensor placement affects the length of the stand still phases since different parts of the boot are stationary during different sequences of the stance. Also, during some movements like running, the boot is stationary for a shorter time and parts like the heel are rarely stationary at all. In this segment we will study how the sensor placement affects the stand still detection performance for different movements.

The first sensor position is by the heel, Figure 3.7a. In a real end user application one idea is to hide the sensor in the thick sole of the heel to protect it, Nilsson et al. (2012). Another option is to put the sensor by the toes since they are involved in almost all movements, Figure 3.7b. Intuitively, the toes should be stationary during pretty much all conceivable standing motions sequences. Since the sensors are getting smaller and smaller, the problem of hiding the sensor in the thinner sole by the toe could soon be solved. A third option is to put the sensor by the shoelaces, Figure 3.7c. This position could be common if the sensor is not integrated with the shoe, but is instead strapped onto the preexisting boot.

The boot type is also important for stand still detection performance. A sturdy boot, Figure 3.7a, rolls over the ground without getting reshaped by the surface while a softer boot, Figure 3.7b, absorbes more of the contact force and alters shape. If the sensor is attached to a sturdy boot, the rolling will result in a sensor that is never really still making stand still phases very hard to detect. A softer boot makes the stand still phases easier to detect but is rarely the boot of choice for firefighters, soldiers etc. The stand still detection performance was evaluated for both boots so the results could be compared.

For the experiments, the sensor was attached using duct tape, Figure 3.7. The sensor is thereby sturdily attached to the boot while it is still very easy to setup experiments. All three sensor positions were evaluated using both boots for walking and running.

The main problem with evaluating even more realistic firefighter movements like crawling is that it is very difficult to obtain reliable ground truth data for such experiments. The experiment has to take place in a high precision positioning lab, for example a VICON[1] lab, but not even such experiments will be completely like actual crawling around furniture and such in an apartment. Since the IMMU data is not really periodic, the stand still phases are not straightforward to pinpoint. The more distinct and repetitive the movement is, the easier it is to differentiate between true and false positives based on the sensor data.

[1] http://www.vicon.com

*(a)* Sturdy hiking boot with the sensor by the heel.

*(b)* Softer winter boot with the sensor by the toe.

*(c)* Sturdy boot with the sensor by the shoe laces.

***Figure 3.7:*** *Sensor positions and available boots for stand still detection experiments.*

#### Experimental Results

A stand still has been assumed detected if the estimated stand still probability is above 90% during a true stand still. If no stand still did occur it is considered a false positive. The 90% threshold is mostly important in the running case. During a walking sequence, the stand still phases are commonly marked out with a stand still probability above 99.9%. In the running sequence the stand still phases are short making it hard to even reach 90% stand still probability.

The preferable result is a combination of high stand still detection ability combined with no or very few false positives. All in all, the best sensor position for stand still detection seems to be a toe mounted sensor due to its enhanced ability to detect stand stills while running.

**Walking** The first movement that was evaluated was walking. This is the most common movement used in localization experiments, Most likely, this is because walking is easy to perform, the data is straightforward to analyze and the stand still phases are long and therefore easy to detect.

The results from the walking sessions are shown in Table 3.1. In general, all stand still phases are detected with very few false positives, no matter the boot or the sensor position.

| **Walking sequence** | | | |
|---|---|---|---|
| Robust boot, 50 steps | Heel | Toe | Lace |
| Stand stills detected | 50 | 50 | 50 |
| False positives | 1 | 0 | 0 |
| Soft boot, 50 steps | Heel | Toe | Lace |
| Stand stills detected | 50 | 50 | 50 |
| False positives | 0 | 0 | 0 |

***Table 3.1:*** *Experimental results of the stand still detection performance of a walking sequence using different sensor positions on two different boots.*

**Running** The second movement is running which is much more difficult. While running, the stand still phases are very short and even then the boot tends to roll quite a lot. Also, when running with sturdy boots the heel is not always even touching the ground.

Table 3.2 shows the stand still detection performance during running. Clearly there are now different results for different settings. It is for once much harder to detect stand stills when a sturdy boot is used than when a soft boot is used. The sturdy boot rolls over the ground making the sensor never really still. Especially the gyro is excited in the angle around the direction pointing straight out to the side. The stand still phases are therefore pretty much non existent when a rigid hiking boot is used. The lace mounted sensors have significant problems detecting any stand stills, even when the softer winter boot is used due to this rolling effect.

One way to improve the results was therefore to remove the problematic gyro dimension. In doing so, the lace mounted rigid boot was able to detect 30 out of the 41 running steps without any false positives, compared to just finding 1 before. Unfortunately, for walking this test statistic instead introduced 19 new false positives.

| **Running sequence** | | | |
|---|---|---|---|
| Robust boot, 41 steps | Heel | Toe | Lace |
| Stand stills detected | 0 | 16 | 1 |
| False positives | 0 | 0 | 0 |
| Soft boot, 41 steps | Heel | Toe | Lace |
| Stand stills detected | 20 | 37 | 0 |
| False positives | 0 | 1 | 0 |

***Table 3.2:*** *Stand still performance during a running sequence using different sensor positions.*

### Discussion

The problem of reliable stand still detection using a foot mounted IMMU is not fully solved yet. In the end the solution will have to be application specific. We have in this section shown that both the type of boot and the position of the sensor can alter the stance detection performance greatly.

These results are not final though. The heel mounted sensor for example will probably experience a different type of movements when actually mounted inside the heel. This would also be true for the toe mounted sensor when the IMMU becomes small enough to hide in the front sole.

Even though the toe mounted sensor is the best for stand still detection, the question does not end there. Different sensor positions excite the sensors differently. The toe mounted sensor for example experiences the biggest angular velocity of the three positions, up to almost 20 rad/s compared to about 10-12 rad/s for the other positions. And this is just during regular walking. The sensor position

thereby puts a higher demand on the sensor quality, requiring that it is accurate for a very large dynamic range.

# 4

# Discussion and Future Work

This chapter ends the first part of the thesis constituted of background theory. The discussion primarily covers the publications which are the second part of the thesis. We will end with suggested directions for future work.

## 4.1 Discussion

Some of the conclusions have been drawn under their respective sections in the publications, but some will be repeated and extended here for completeness.

### 4.1.1 Indoor Localization

Our work in indoor localization for professional users covers the problem of positioning in small or medium sized venues of which there is no prior information available. The approach taken was to extend and improve the foot mounted IMMU based dead reckoning systems suggested in the literature. It is a simple yet fundamental system that should be a part of a final positioning system.

The main shortcoming of such systems is the drift in position estimate. This is primarily caused by a drift in heading, not in distance traveled. The stand still detection and utilization in the filtering has significantly reduced the latter. To reduce the drift in heading, approaches like adding sensors, for example cameras, have been taken. Also extensive system calibration has been used to reduce the drift but this is not really a solution to the problem, only a postponement. Our feeling is that one can do more with just one IMMU. Before additional sensors were added, had the limit for positioning using only one IMMU really been reached?

Our primary goal hence became solving the heading estimation problem. If one

could have both a reliable heading estimate and movement length estimate, the final position estimate should be quite good. To do this we decided to use the sensors already available in the IMMU: the magnetometer and the gyro.

Two heading estimation methods have been presented. The first one presented in Paper A is a Kalman filter bank based filtering method. Two modes are estimated, disturbed and undisturbed, deciding which magnetometer measurements can be used for heading estimation and which should be discarded. The final result is a filtering estimate that is significantly better than the raw magnetometer measurements in a vast majority of all experiments.

The second system presented in Paper B is an optimization based approach. The estimation is done in batch form over the entire dataset producing a smoothing estimate for the heading of the complete experiment. The strength of the system is its ability to handle significant disturbances while maintaining the smoothness of the gyro integration trajectory. A weakness is the somewhat ad hoc approach and that its basic assumptions of the gyro quality is perhaps not valid for all sensors.

The heading estimation approaches have been merged with a dead reckoning system. The system is based on a probabilistic stand still detection framework described in Section 3.3.1. The HMM based stand still detection produces a probability of standstill that experiments verify is very reliable. The best stand still detection ability has turned out to be achieved when the sensor is mounted near the toe. It has also been shown that a more rigid boot makes stand still detection more difficult since it has a tendency to roll over the ground. The stand still detection system cannot handle running well today, but it could probably be solved by also incorporating the step frequency that will be present in such a scenario.

To include the heading estimation system in the positioning system, it is crucial that it is only used once per stand still phase. If the heading is estimated multiple times per stand still, basically the exact same measurement of magnetometer heading and system orientation will be used by the heading estimation system over and over again. This will ruin the heading estimation since magnetic disturbances become extremely hard to detect.

Including the magnetic heading measurements in the stand still update is not a good approach. It can be done using for example the IMM based heading estimation framework in Paper A. Unfortunately, the resulting estimated position trajectory becomes very uncertain. Since a small change in orientation will cause the position estimates to jump significantly, the position trajectory becomes very shaky. To make the trajectory smooth, all previous position estimates must be adjusted using smoothing.

The approach we suggest is to use accelerometer and gyro to produce estimates of step to step movements using stand still detection and measurement integration. These estimates are then reduced to estimates in 2D, $\hat{x}_{0:N}$, $\hat{y}_{0:N}$ and heading $\hat{\psi}_{0:N}$. This vector of gyro based heading estimates are then adjusted to match a vector of magnetic heading measurements to remove the drift. These new head-

ing estimates $\hat{\hat{\psi}}_{0:N}$ are finally used to produce new position estimates $\hat{\hat{x}}_{0:N}, \hat{\hat{y}}_{0:N}$. The experimental results show that the estimated user trajectory improves significantly when the magnetometer is included in the estimation.

### 4.1.2 RADAR SLAM

Naval vessels today rely heavily on the GNSS system. Not only is it used for positioning, but also for its exact time estimates. It is not unusual that a large number of control and decision systems are receiving information from the GNSS system. That means that if GNSS fails, a number of crucial systems on the ship will fail simultaneously. In controlled experiments, GNSS jamming has turned out to cause complete chaos on the bridge. This makes naval vessels very vulnerable to GNSS disturbances, intentional or unintentional, in critical situations.

To reduce the impact of GNSS outage, we have designed a positioning system that will produce accurate estimates of velocity, heading and position of the vessel throughout the outage, using only the vessel's own RADAR sensor. This system is presented in Paper D. It uses the reflections from the surrounding islands and stable landmarks are tracked over time. By studying how the islands move relative to the vessel, its position can be estimated.

A strength of the system is its simplicity. All commercial vessels are already equipped with a RADAR sensor why no new expensive equipment is needed. One weakness of the system is of course that if there are no islands available to track, no position estimate can be produced. This is unfortunate, but if one experiences GNSS outage in the middle of the ocean, the situation is not as critical as if one is in an archipelago or harbor, why no backup system is really needed in the first place. So the system will only work if the vessel is surrounded by islands or a coastline, but this is also the only scenario in which it would actually be needed.

### 4.1.3 Underwater Sensor Positioning

A passive underwater sensor localization scheme is presented in Paper E. It uses triaxial magnetometers and a friendly vessel with known magnetic characteristics to determine the sensors positions.

Simulations indicate that if the vessel is equipped with GNSS, a majority of all sensors can be positioned accurately. The simulations also indicate that our positioning scheme is quite insensitive to minor errors in sensor orientation and magnetic signature, when GNSS is used to position the vessel throughout the trajectory.

The strength of the system is that it solves a challenging positioning problem using a very cheap method. The only thing needed is a friendly vessel that travels through the area. The weakness is that it requires that the magnetic signature of the vessel is quite well known in advance. That means that a different system is needed to estimate the signature if it is not actively produced on board. There is also quite a lot of synchronization needed to ensure that all measurements can be used jointly. The measurements from all sensors and the vessel must be

synchronized for this to work in practice. In the end quite a lot of measurement management will probably be needed in the estimation system.

## 4.2 Future Work

The indoor positioning system described in this thesis is performing well but is quite limited. For example, it does not provide any information about the building layout to the operational management. The next step now that the dead reckoning is stable, is to incorporate more sensors in the system to produce more reliable and informative localization.

A basic extension is to use multiple IMMUs. When only one foot mounted IMMU is used, some motions like crawling is hard to track. While crawling, the feet have a tendency of wobbling why correctly detecting stand stills become difficult. A knee mounted IMMU would solve this since the knee is often stationary while crawling. The system would then switch between using the foot mounted sensor and the knee mounted sensor depending on what the data suggests the user is doing.

Joint positioning is another option. If multiple users are in the building, measuring the distances between the users will restrict the overall error growth. One way of solving this would be to communicate the estimated user positions and their relative distances to each other to an outside system that use this to produce a joint estimate of all positions.

Other sensors like cameras or infrared cameras can be used. Stable landmarks in the environment can be tracked from image to image, giving estimates of how the user is moving. The camera can be used with an IMMU to produce even more robust position estimates. If an area is revisited, camera images can be used to detect this, a so called loop closure. The position estimate can then be updated with this information. If multiple cameras equipped users are in the building together, one user can get a loop closure from another user's images. That means that one user does not have to revisit a location for it to be a loop closure. If the images from one user can be connected to the images from a second user, their estimated positions can be connected. This could be useful in the scenario where police officers or soldiers search through a building.

The trucks used to get to the scene can be equipped with transponders. These signals could be picked up by sensors worn by the user to aid positioning. Unfortunately, signal multipathing might be an issue in such a system. Such a system would probably also have deteriorating performance with building size. It might be working well in a residential house but not in a skyscraper.

The underwater positioning system needs experimental validation but this requires expensive underwater sensors with cables on the sea floor and a large enough steel vessel to be used as target. Since we currently have neither sensors nor vessel, this has been put on hold.

# Appendix

# A

# Quaternion Properties

Quaternions were discovered by Hamilton as an extension for the imaginary numbers into three dimensions, Hamilton (1844). Later, the unit quaternion started to be used for angle representation in rotations providing singularity free rotations. For a thorough description of quaternions see Kuipers (1999).

## A.1 Operations and Properties

A quaternion is a four-tuple of real numbers denoted by $q = (q_0, q_1, q_2, q_3)$. Alternatively it can be described as consisting of a scalar part $q_0$ and the vector $\mathbf{q}$.

$$q = \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix} \tag{A.1}$$

Quaternion multiplication is denoted as $\odot$ and defined as

$$p \odot q = \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} \odot \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} p_0 q_0 - \mathbf{p} \cdot \mathbf{q} \\ p_0 \mathbf{q} + q_0 \mathbf{p} + \mathbf{p} \times \mathbf{q} \end{pmatrix}. \tag{A.2}$$

Some quaternions properties are:

$$p \odot q \neq q \odot p \tag{A.3}$$

$$\text{norm}(q) = \sqrt{\sum_{i=0}^{3} q_i^2} \tag{A.4}$$

$$q^{-1} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix}^{-1} = \begin{pmatrix} q_0 \\ -\mathbf{q} \end{pmatrix} \tag{A.5}$$

The unit quaternion used for rotation operations also fulfills $\text{norm}(q) = 1$.

## A.2 Describing a Rotation using Quaternions

The quaternion

$$q = \begin{pmatrix} \cos\delta \\ \sin\delta\mathbf{n} \end{pmatrix} \tag{A.6}$$

describes a rotation around the vector $\mathbf{n}$ with the angle $2\delta$. There are two ways of depicting a rotation: either the coordinate frame is rotated and the vector is fixed or the vector is rotated and the coordinate frame is fixed. The difference is in the sign of the rotation. In this work the vector will be assumed constant and the coordinate system is rotated.

A rotation of a vector $v$ around $\mathbf{n}$ can be written as

$$u = q^{-1} \odot v \odot q \tag{A.7}$$

where

$$v = \begin{pmatrix} 0 \\ \mathbf{v} \end{pmatrix} \quad \text{and} \quad u = \begin{pmatrix} 0 \\ \mathbf{u} \end{pmatrix}. \tag{A.8}$$

This is the assumed standard rotation in this work. The resulting rotation is

$$v = \begin{pmatrix} \mathbf{q} \cdot \mathbf{u} q_0 - (q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) \cdot \mathbf{q} \\ (\mathbf{q} \cdot \mathbf{u})\mathbf{q} + q_0(q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) + (q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) \times \mathbf{q} \end{pmatrix} \tag{A.9}$$

which simplifies to

$$v = \begin{pmatrix} 0 \\ 2(\mathbf{q} \cdot \mathbf{u})\mathbf{q} + (q_0^2 - \mathbf{q} \cdot \mathbf{q})\mathbf{u} - 2q_0\mathbf{q} \times \mathbf{u} \end{pmatrix}. \tag{A.10}$$

## A.3 Rotation Matrix

The quaternion rotation (A.10) can be rewritten as a matrix multiplication

$$\mathbf{v} = R(q)\mathbf{u} \tag{A.11}$$

where

$$R(q) = \begin{pmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2 & 2(q_1 q_2 + q_0 q_3) & 2(q_1 q_3 - q_0 q_2) \\ 2(q_1 q_2 - q_0 q_3) & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2(q_2 q_3 + q_0 q_1) \\ 2(q_1 q_3 + q_0 q_2) & 2(q_2 q_3 - q_0 q_1) & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{pmatrix}. \tag{A.12}$$

See Kuipers (1999) for details.

## A.4 Quaternion Dynamics

In the case of the quaternions describing a rotation between a global coordinate system and a local one attached to a moving sensor, the description of the quaternions will contain some dynamics. The full derivation of the quaternion dynamics can be studied in Kuipers (1999) but parts will be recited here.

Let the quaternion $q_{lg}$ represent the rotation of the local coordinate system in the global one. The angular velocity of the sensor unit in the local coordinate system is $\omega_{lg}^{l}$ which can be written as

$$\omega_{lg}^{l} = \begin{pmatrix} 0 \\ \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} = \begin{pmatrix} 0 \\ \boldsymbol{\omega} \end{pmatrix} \tag{A.13}$$

while $q_{lg}$ is

$$q_{lg} = \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix}. \tag{A.14}$$

The quaternion derivative can now be written as

$$\begin{aligned} \dot{q}_{lg} &= \frac{1}{2} q_{lg} \odot \omega_{lg}^{l} = \frac{1}{2} \begin{pmatrix} -\mathbf{q} \cdot \boldsymbol{\omega} \\ q_0 \boldsymbol{\omega} + \mathbf{q} \times \boldsymbol{\omega} \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} -(q_1 \omega_x + q_2 \omega_y + q_3 \omega_z) \\ q_0 \boldsymbol{\omega} - \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix} \end{pmatrix} \\ &= \underbrace{\frac{1}{2} \begin{pmatrix} 0 & -\omega_x & -\omega_y & -\omega_z \\ \omega_x & 0 & \omega_z & -\omega_y \\ \omega_y & -\omega_z & 0 & \omega_x \\ \omega_z & \omega_y & -\omega_x & 0 \end{pmatrix}}_{S(\omega_{lg}^{l})} \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \underbrace{\frac{1}{2} \begin{pmatrix} -q_1 & -q_2 & -q_3 \\ q_0 & -q_3 & q_2 \\ q_3 & q_0 & -q_1 \\ -q_2 & q_1 & q_0 \end{pmatrix}}_{S'(q_{lg})} \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} \end{aligned}$$

If the angular velocity $\omega_k$ is assumed constant over the sampling interval the noise free discrete time model is

$$q_{k+1} = e^{\frac{1}{2}S(\omega_k)T} q_k \tag{A.15}$$

where a Taylor series expansion gives

$$\begin{aligned} e^{\frac{1}{2}S(\omega_k)T} &= \sum_{n=0}^{\infty} \frac{\left(\frac{1}{2}S(\omega_k)T\right)^n}{n!} \\ &= \sum_{n=0}^{\infty} \left(\frac{T}{2}\right)^n \frac{1}{n!} S(\omega_k)^n. \end{aligned} \tag{A.16}$$

If the sampling time $T$ is short, the expansion can be approximated with the first two terms

$$e^{\frac{1}{2}S(\omega_k)T} \approx I + \frac{T}{2}S(\omega_k) \tag{A.17}$$

giving the discrete time model

$$\begin{aligned} q_{k+1} &= q_k + \frac{T}{2}S(\omega_k)q_k \\ &= q_k + \frac{T}{2}S'(q_k)\omega_k. \end{aligned} \tag{A.18}$$

If the angular velocity noise term $\nu_\omega$ is included the discrete model becomes

$$q_{k+1} = q_k + \frac{T}{2}S'(q_k)\omega_k + \frac{T}{2}S'(q_k)\nu_{\omega,k}. \tag{A.19}$$

# Bibliography


K. Abdulrahim, C. Hide, T. Moore, and C. Hill. Aiding low cost inertial navigation with building heading for pedestrian navigation. *Journal of Navigation*, 64: 219–4233, 2011.

P. Aggarwal, D. Thomas, L. Ojeda, and J. Borenstein. Map matching and heuristic elimination of gyro drift for personal navigation systems in GPS-denied conditions. *Measurement Science and Technology*, 22(2):025205, 2011.

M. Angermann and P. Robertson. FootSLAM: Pedestrian simultaneous localization and mapping without extroceptive sensors – hitchiking on human perception and cognition. *Proceedings of the IEEE*, 100(13):1840–1848, 2012.

P. Bahl and V.N Padmanabhan. RADAR: an in-building rf-based user location and tracking system. In *Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies INFOCOM*, 2000.

T. Bailey and H. Durrant-Whyte. Simultaneous localization and mapping (SLAM): part II. *Robotics & Automation Magazine, IEEE*, 13(3):108 – 117, 2006.

Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. Wiley, New York, 2001.

S. Beauregard. Omnidirectional pedestrian navigation for first responders. In *Proc. of the 4th Workshop on Positioning, Navigation and Communication, WPNC07*, Hannover, Germany, 2007.

S. Beauregard, Widyawan, and M. Klepal. Indoor PDR performance enhancement using minimal map information and particle filters. In *IEEE/ION Position, Location and Navigation Symposium*, 2008.

Ö. Bebek, M.A. Suster, S. Rajgopal, M.J. Fu, X. Huang, M.C. Çavuşoğlu, D.J. Young, M. Mehregany, A.J. van den Bogert, and C.H. Mastrangelo. Personal navigation via high-resulution gait-corrected inertial measurement units. *IEEE Transactions on Instrumentation and Measurement*, 59(11):3018–3027, 2010.

J. Borenstein and L. Ojeda. Heuristic drift elimination for personnel tracking systems. *Journal of Navigation*, 63(3):591 – 606, 2010.

J. Callmer. Topics in localization and mapping. Linköping University, Sweden, 2011. Licentiate Thesis No 1489.

J. Callmer, K. Granström, J. Nieto, and F. Ramos. Tree of words for visual loop closure detection in urban SLAM. In *Proceedings of the 2008 Australasian Conference on Robotics and Automation (ACRA)*, 2008.

J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010a.

J. Callmer, D. Törnqvist, and F. Gustafsson. Probabilistic stand still detection using foot mounted IMU. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010b.

J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors. *IEEE Transactions on Signal Processing*, 2013a. Submitted.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors using convex optimization. In *International Conference on Information Fusion*, 2013b. Submitted.

J. Callmer, D. Törnqvist, and F. Gustafsson. An inertial navigation framework for indoor positioning with robust heading. *IEEE Transactions on Instrumentation and Measurement*, 2013c. Submitted.

B.J. Dil and P.J.M. Havinga. RSS-based localization with different antenna orientations. In *Australasian Telecommunication Networks and Applications Conference*, 2010.

H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping (SLAM): part I. *Robotics & Automation Magazine, IEEE*, 13(2):99 – 110, 2006.

The Economist. No jam tomorrow. *The Economist Technology Quarterly*, 398 (8724):20–21, 2011.

The Economist. Finding the way inside. *The Economist Technology Quarterly Q4*, 2012.

E. Erkstam and E. Tjernqvist. Performance enhancement of bearing navigation to known radio beacons. Master's thesis, Linköping University, 2012.

R. Feliz, E. Zalama, and J. G. Garcia-Bermejo. Pedestrian tracking using inertial sensors. *Journal of Physical Agents*, 3(1):35–43, 2009.

E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

S. Godha, G. Lachapelle, and M. E. Cannon. Integrated GPS/INS system for pedestrian navigation in signal degraded environment. In *Proc. of ION GNSS*, 2006.

K. Granström, J. Callmer, F. Ramos, and J. Nieto. Learning to detect loop closure from range data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.

A. Grant, P. Williams, N. Ward, and S. Basker. GPS jamming and the impact on maritime navigation. *Journal of Navigation*, 62(2):173–187, 2009.

S. Grzonka, F. Dijoux, A. Karwath, and W. Burgard. Mapping indoor environments based on human activity. In *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, Anchorage, Alaska, 2010.

Y. Gu, A. Lo, and I. Niemegeers. A survey of indoor positioning systems for wireless personal networks. *IEEE COMMUNICATIONS SURVEYS and TUTORIALS*, 11(1):13–32, 2009.

F. Gustafsson. *Statistical Sensor Fusion*. Studentlitteratur, 2010.

S.W.R. Hamilton. On quaternions; or on a new system of imaginaries in algebra. *Philosophical Magazine*, xxv:10–13, 1844.

G. Hanning, N. Forslöw, P.-E. Forssén, E. Ringaby, D. Törnqvist, and J. Callmer. Stabilizing cell phone video using inertial measurement sensors. In *In Proceedings of the IEEE International Workshop on Mobile Vision (IWMV11)*, 2011.

V. Honkavirta, T. Perala, S. Ali-Loytty, and R. Piché. A comparative survey of WLAN location fingerprinting methods. In *Workshop on Positioning, Navigation and Communication (WPNC 2009)*, 2009.

A.H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, 1970.

A.R. Jiménez, F. Seco, J.C. Prieto, and J. Guevara. Indoor pedestrian navigation using an ins/ekf framework for yaw drift reduction and a foot-mounted imu. In *Positioning Navigation and Communication (WPNC), 2010 7th Workshop on*, pages 135 –143, march 2010a.

A.R. Jiménez, F. Seco, J.C. Prieto, and J. Guevara. Indoor pedestrian navigation using an ins/ekf framework for yaw drift reduction and a foot-mounted imu. In *7th Workshop on Positioning, Navigation and Communication, WPNC*, 2010b.

S. Kaisera, M. Khidera, and P. Robertson. A pedestrian navigation system using a map-based angular motion model for indoor and outdoor environments. *Journal of Location Based Services*, 7(1), 2013.

R.E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME J Basic Engr.*, 82:35–45, 1960.

B. Krach and P. Robertson. Cascaded estimation architecture for integration of foot-mounted inertial sensors. In *Proc of the IEEE/ION Position, Location and Navigation Symposium*, 2008.

J.B. Kuipers. *Quaternions and Rotation Sequences*. Princeton University Press, 1999.

B. Li, J. Salter, A.G. Dempster, and C. Rizos. Indoor positioning techniques based on wireless LAN. Technical report, School of Surveying and Spatial Information Systems, UNSW, Sydney, Australia, 2006.

F. Lindsten, J. Callmer, H. Ohlsson, D. Törnqvist, T. B. Schön, and F. Gustafsson. Geo-referencing for UAV navigation using environmental classification. In *Proceedings of 2010 International Conference on Robotics and Automation (ICRA)*, 2010.

C. Lundquist, P. Skoglar, F. Gustafsson, D. Törnqvist, and J. Callmer. Method and device for indoor positioning. US Patent Application 20120203453, August 8, 2012.

K. Muthukrishnan, M.E. Lijding, and P.J.M. Havinga. Towards smart surroundings: Enabling techniques and technologies for localization. In *Proceedings of the International Workshop on Location-and Context-Awareness*, 2005.

J.-O. Nilsson, I. Skog, P. Händel, and K. Hari. Foot-mounted INS for everybody - an open-source embedded implementation. In *Position Location and Navigation Symposium (PLANS), 2012 IEEE/ION*, pages 140–145, 2012.

L. Ojeda and J. Borenstein. Non-GPS navigation for security personnel and first responders. *Journal of Navigation*, 60(3):391–407, 2007.

J. Rantakokko, J. Rydell, P. Strömbäck, P. Händel, J. Callmer, D. Törnqvist, F. Gustafsson, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: multisensor systems and cooperative localization. *Wireless Communications, IEEE*, 18(2):10–18, 2011.

V. Renaudin, O. Yalak, P. Tomé, and B. Merminod. Indoor navigation of emergency agents. *European Journal of Navigation*, 5(3), 2007.

P. Robertson, M. Angermann, and B. Krach. Simultaneous localization and mapping for pedestrians using only foot-mounted inertial sensors. In *Proceedings of the 11th international conference on Ubiquitous computing*, Ubicomp '09, pages 93–96, 2009.

T. Roos, P. Myllymäki, H. Tirri, P. Misikangas, and J. Sievänen. A probabilistic approach to WLAN user location estimation. *International Journal of Wireless Information Networks*, 9(3), 2002.

J. Rydell and T. Andersson. Camera-based navigation, mapping and 3d model generation for indoor environments. In *Proceedings of the 2010 International Technical Meeting of The Institute of Navigation*, 2010.

J. Seitz, T. Vaupel, J. Jahn, S. Meyer, J.G. Boronat, and J. Thielecke. A hidden markov model for urban navigation based on fingerprinting and pedestrian dead reckoning. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

I. Skog. *Low-cost Navigation Systems*. PhD thesis, KTH, Stockholm, Sweden, 2009.

I. Skog, P. Händel, J.-O. Nilsson, and J. Rantakokko. Zero-velocity detection – an algorithm evaluation. *IEEE Transactions on Biomedical Engineering*, 57, 2010.

N. Wahlström, J. Callmer, and F. Gustafsson. Magnetometers for tracking metallic targets. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

N. Wahlström, J. Callmer, and F. Gustafsson. Single target tracking using vector magnetometers. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.

H. Wang, H. Lenz, A. Szabo, J. Bamberger, and U.D. Hanebeck. WLAN-based pedestrian tracking using particle filters and low-cost MEMS sensors. In *4th workshop in positioning, navigation and communication (WPNC)*, 2007.

Widyawan, M. Klepal, and D. Pesch. A bayesian approach for RF-based indoor localisation. In *4th International Symposium on Wireless Communication Systems ISWCS*, 2007.

Widyawan, M. Klepal, and S. Beauregard. A backtracking particle filter for fusing building plans with PDR displacement estimates. In *Positioning, Navigation and Communication, 2008. WPNC 2008. 5th Workshop on*, pages 207 –212, march 2008.

O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing (UbiComp)*, 2008.

O. Woodman and R. Harle. Rf-based initialisation for inertial pedestrian tracking. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

H. Wymeersch and M.Z. Win J. Lien. Cooperative localization in wireless networks. *Proceedings of the IEEE*, 97(2):427–450, 2009.

M. Youssef and A. Agrawala. The horus WLAN location determination system. In *Proceedings of the 3rd international conference on Mobile systems, applications and services (MobiSys)*, 2005.

D. Zachariah and M. Jansson. Fusing visual tags and inertial information for indoor navigation. In *IEEE/ION Position Location and Navigation Symposium (PLANS)*, 2012.

# Part II

# Publications

# Paper A

# Robust Heading Estimation Indoors

*Authors:* Jonas Callmer, David Törnqvist and Fredrik Gustafsson

*Edited version of the paper:*

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors. *IEEE Transactions on Signal Processing*, 2013a. Submitted.

*Preliminary version:*

Technical Report LiTH-ISY-R-3061, Dept. of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden.

# Robust Heading Estimation Indoors

Jonas Callmer, David Törnqvist and Fredrik Gustafsson

Dept. of Electrical Engineering,
Linköping University,
SE–581 83 Linköping, Sweden
callmer@isy.liu.se

**Abstract**

Indoor positioning in unknown environments is crucial for rescue personnel and future infotainment systems. Dead-reckoning inertial sensor data gives accurate estimate of distance, for instance using zero velocity updates, while the heading estimation problem is inherently more difficult due to the large degree of magnetic disturbances indoors. We propose a Kalman filter bank approach based on supporting a magnetic compass with gyroscope turn rate information, where a hidden Markov model is used to model the presence of magnetic disturbances. In parallel, we suggest to run a robust heading estimation system based on data from a sliding window. The robust estimate is used to detect filter divergence, and to restart the filter when needed. The underlying assumptions and the heading estimation performance are supported in field trials using more than 500 data sets from more than 50 venues in 5 continents.

## 1 Introduction

We consider the problem of personal navigation in unknown indoor environments, which is a typical situation for instance for rescue personnel, police officers or soldiers.

The assumptions for personal navigation indoors in our work are as follows:

**A1** No radio map.

**A2** No building map or building structure available.

**A3** The only available information comes from low-price MEMS sensors in an inertial magnetic measurement unit (IMMU) comprising accelerometer, gyroscope and magnetometers, all being three-dimensional.

**A4** The earth magnetic field is disturbed by hard and soft iron such as steel structures, electrical wiring, etc.

The first assumption excludes satellite navigation and also all kind of fingerprinting or triangulation from radio beacons Pahlavan et al. (2010). With a building map, we would be in a favorable situation where the movements are restricted by known walls and obstacles, which is sufficient extra information to support dead-reckoning Gusenbauer et al. (2010); Rantakokko et al. (2011). However, maps are seldom available, hence assumption two. Further, the building does not need to be restricted to corridors with 90 degrees separation, or any other *a priori* known structure.

Accurate estimation of distance can be achieved with for instance a foot mounted IMMU, where zero velocity detection (ZUPT) is used to correct built up errors in velocity, which decreases the drift in distance from cubic to linear in time Foxlin (2005).

In a dead reckoning system for indoor positioning based on a foot mounted IMMU, position is updated at time $k$ when a new step $s_k$ is available. The position $p_k$ is updated as

$$\begin{pmatrix} p_k^x \\ p_k^y \end{pmatrix} = \begin{pmatrix} p_{k-1}^x \\ p_{k-1}^y \end{pmatrix} + s_k \begin{pmatrix} \cos(\psi_k) \\ \sin(\psi_k) \end{pmatrix} \tag{1}$$

where $\psi_k$ is the heading in which the step was taken. Clearly, if $\psi_k$ is wrong, $p_k$ will be wrong, even if $s_k$ is correct. Therefore, to be able to accurately position personell indoors using an unsupported dead reckoning system, an accurate estimate of $\psi_k$ is needed.

In the end, the most challenging yet crucial problem in indoor positioning in unknown environments is the estimation of heading.

Since no external support such as radio environment or maps of floor layout can be used (assumption 1 and 2), we are left with only low grade magnetometer and gyro measurements. The magnetometer measurements would, if the magnetic field is undisturbed, give the direction of north, but the magnetic field is often suffering from major disturbances indoors that are typically location dependent. The fourth assumption therefore excludes just using the magnetometers to measure $\psi_k$ The gyro can only measure the change in user orientation, but it does it quite well. Unfortunately, most often the sensor, especially a low grade one, has slight inaccuracies which over time introduce an ever growing error in the orientation estimate.

However, together the magnetometer and gyroscopes have complementary properties. Figure 1 illustrates the situation plotting the magnetic heading measurements and the ground truth for a short experiment. Also shown is the heading estimate acquired by summing up the angular velocity measurements from the gyro. It is clear that without external support an estimate using only gyro measurements just drifts away. Our key idea is to support the gyroscope and learn its bias and drift during occasional segments with reliable magnetometer data. To do this we also need to estimate when the magnetometer data is reliable and when it is not.

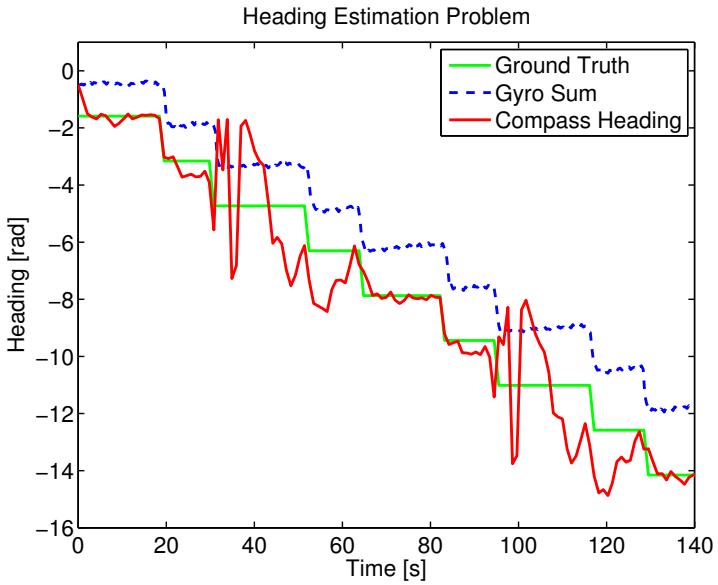


***Figure 1:*** *How do we fuse the disturbed magnetometer signal (solid red) with the biased gyro measurements (dashed blue) to reproduce the ground truth (solid green)?*

To estimate $\psi$, a Kalman filter bank is used where one filter assumes that the magnetometer signal is disturbed while the other one assumes that it is not. During the disturbed data phases, the estimate is forced to rely on gyro measurements only. To reduce the drift in estimate during the disturbances, gyro error correction terms are estimated using the disturbance free sections of the magnetometer measurements.

The outline is as follows. Section 2 surveys related work on how to model magnetic disturbances and remedies for indoor positioning. Section 3 investigates the magnetometer signal in more detail. Section 4 provides the mathematical framework and our experience of sensor data from field tests, while Section 5 describes the models and optimal filter we propose. Section 6 proposes a robust, but not optimal, filter based on data over a sliding window only. This is used to monitor the performance of the optimal filter, and restart it when necessary. There are many implementation aspects of our general framework, and Section 7 describes one particular method used in the experiments. Section 8 first analyzes two data sets in detail, and then summarized more than 500 field tests. Section 9 concludes the paper. Note that throughout the work, the terms yaw and heading will be used interchangeably.

# 2 Related Work on Magnetic Disturbances Indoors

The related work covers four fields. First, there are studies of the nature of the magnetic disturbances. Second, some work attempts to reduce the impact of the magnetometer signal disturbances by reducing the influence of the magnetometer signal in orientation estimation. Third, we have the approaches to detect and sometimes estimate the magnetic disturbances to reduce their impact. Finally, there are indoor positioning approaches based on magnetometer free methods to solve the heading estimation problem.

## 2.1 Disturbance Studies

The appearance of the disturbances on the magnetometer measurements are thoroughly described in Bachmann et al. (2007) and de Vries et al. (2009). In the former it is shown how different ferrous objects and electrically powered devices affect the earth magnetic field and also how the disturbances depend on the distance to the interfering object. In the latter the magnetic fields of a lab have been mapped extensively. The results confirm that the disturbances decay rapidly with distance.

## 2.2 Influence Reduction

In Suh et al. (2012); Yun et al. (2008) they estimate 3D orientation where the magnetometer is seen as problematic. They present estimation methods based on an accelerometer and a magnetometer that only allow the magnetometer measurements to influence the yaw estimate, not the pitch and roll estimates.

## 2.3 Disturbance Detection

The most common approach to handling magnetic disturbances is to study the signal norm and/or the dip angle Sabatini (2006); Lee and Park (2009); Harada et al. (2004); Sabatini (2011). If either deviate from a specific interval, the magnetic signal is deemed unreliable and is subsequently discarded.

A kind of soft decision of reliability is proposed in Roetenberg et al. (2005). In the filter, the magnetometer measurement noise is modeled as a function of the signal norm and the dip angle. The disturbance is also estimated using a random walk model. The system is intended for human body segment orientation estimation and not for indoor positioning, why the experiments are much more limited in time and especially in space.

A binary mode estimation approach is presented by Kang and Park (2010). They estimate attitude using accelerometer and magnetometer data, where the validity of the data is estimated using a Hidden Markov Model that studies the magnetometer signal norm. The resulting mode affects the parameter settings of a Kalman filter that estimates the attitude.

### 2.4 Alternative Yaw Estimation Approaches

Most often the magnetometer signal has been discarded as unreliable indoors and simply ignored Stirling et al. (2003); Borenstein and Ojeda (2010). Alternative ways to estimate heading have therefore been suggested Bachmann et al. (2012); Borenstein et al. (2009); Borenstein and Ojeda (2010).

Bachmann et al. (2012) presents a positioning system where the heading is estimated using only gyros. They measure the positioning error acquired during an offline experiment, which is used to calculate a minor heading drift correction term once and for all. Borenstein et al. (2009); Borenstein and Ojeda (2010) make the basic assumption that the user is most of the time walking in a specific direction. In Borenstein et al. (2009) the angular velocity in yaw is pushed towards zero to make the paths more straight. In Borenstein and Ojeda (2010) they instead push the heading into 4-8 predefined directions. It works fine only if the user is actually walking in straight lines most of the time.

## 3 Magnetometer Signal Evaluation

We will now study the magnetometer signal in more detail. What does the data from a real experiment look like and what conclusions can be drawn from that?

First, in Sabatini (2006); Lee and Park (2009); Harada et al. (2004); Sabatini (2011) they study the magnetometer signal norm and/or the dip angle to find trustworthy magnetometer data. In theory it should be a safe method, but in practice we have found it not to be so and we will now show why.

Figure 2 shows data from an outdoor dataset with measured yaw angle with ground truth, the dip angle and the magnetic signal norm. Outdoors the magnetic field should be undisturbed but even here there are significant fluctuations in the data. The dip angle should be 61° and the mean angle in the experiment is 57° which is close, but it fluctuates between 45° and 70°. The magnetic norm is more steady with a mean of 51, and a min and a max of 45 and 56, respectively.

Indoors, Figure 3, the dip angle and the norm are not consistent for disturbed and undisturbed data. Where the yaw angle is good, the dip angle can be anything from 11° to 62°, while the magnetic norm covers a span from 32 to 75. This is not so different from the disturbed periods where the dip angle is between −3° and 82° and the magnetic signal norm is from 15 to 54.

There is apparently a significant overlap in dip angle and magnetic signal norm between the disturbed and undisturbed data segments. Our conclusion is therefore that studying the magnetic signal norm and the dip angle is not a safe method to detect trustworthy magnetic yaw measurements.

Figure 3 also shows us some other features of the magnetometer signal. Most importantly: the disturbances come in bursts. That means that a disturbed measurement is most likely followed by another disturbed measurement and the same is

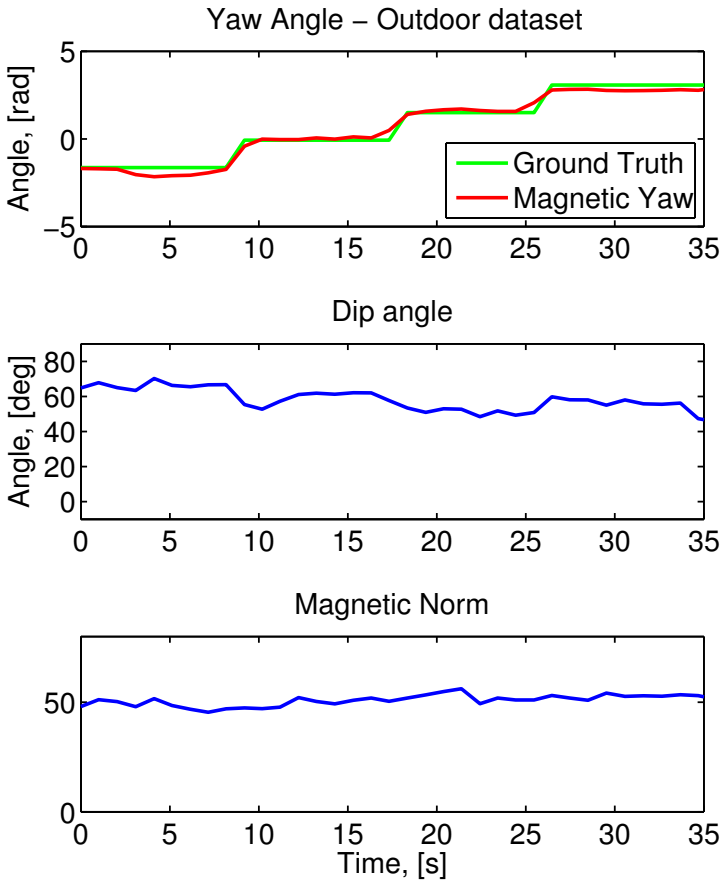


***Figure 2:*** *Yaw measurements with ground truth, dip angle and magnetic signal norm in an outdoor experiment.*

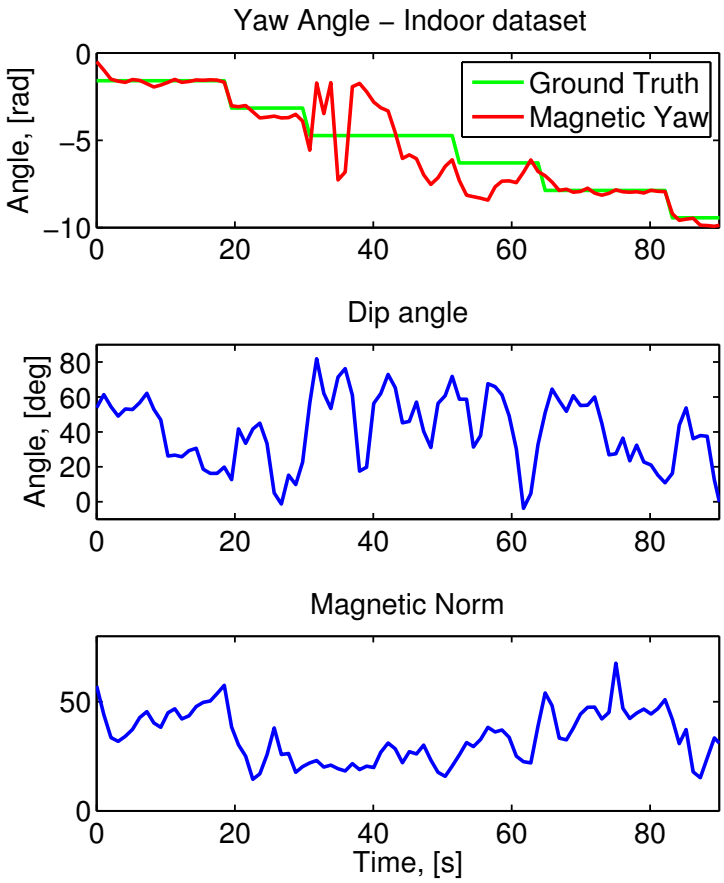


***Figure 3:*** *Yaw measurements with ground truth, dip angle and magnetic signal norm in an indoor experiment with significant disturbances. Note how the dip angle and the magnetic signal norm fluctuate even when the yaw measurements are good.*

true for undisturbed measurement. Producing a good yaw estimate therefore requires a more advanced filtering scheme than just discarding the occasional poor measurement. It also shows us that the disturbances can create heading measurements of any value, sometimes even close to the correct one.

## 4 Principles of Yaw Estimation

The sensors measure at time $k$ three-dimensional acceleration $\mathbf{y}_k^a$, angular rate acceleration $\mathbf{y}_k^\omega$, and magnetic field $\mathbf{y}_k^m$, respectively. The goal is to estimate the yaw angle $\psi_k$.

Suppose that the IMMU is kept horizontally. A first principle for heading estimation is based on integrating angular rates. The third component of the gyroscope gives a measure of yaw rate,

$$y_k^{\dot{\psi}} = (0,\ 0,\ 1)\mathbf{y}_k^\omega, \tag{2}$$

so the yaw angle can be determined from

$$\hat{\psi}_k^1 = \psi_0 + \sum_{l=1}^{k}(0,\ 0,\ 1)\mathbf{y}_l^\omega. \tag{3}$$

A second principle comes from the property that the earth magnetic field is directed to north ($Y$ in world coordinates), while gravity is downwards ($-Z$). The accelerometer though measure the normal force, why $\mathbf{y}_k^a$ is not a measurement of down but of up.

Let us now assume that the sensor is at rest and that it is kept horizontally with yaw angle $\psi$. The cross product

$$\mathbf{z}_k^E = \mathbf{y}_k^m \times \mathbf{y}_k^a = R(\psi)\begin{pmatrix} 0 \\ \cos(\delta)m \\ -\sin(\delta)m \end{pmatrix} \times \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix} \tag{4}$$

$$= \begin{pmatrix} gm\cos(\delta)\cos\psi \\ gm\cos(\delta)\sin\psi \\ 0 \end{pmatrix} \tag{5}$$

gives the vector pointing east ($X$) in the sensor coordinate system. Here, $m$ denotes the strength and $\delta$ is the dip angle of the magnetic field, respectively, and $R(\psi)$ is the rotation matrix denoting a rotation in yaw.

The yaw angle $\psi$ can then be determined as

$$\hat{\psi}_k^2 = \arctan(\mathbf{z}(2)_k^E/\mathbf{z}(1)_k^E). \tag{6}$$

The problems in practice have the following causes:

- There is noise on all sensor signals, assumption 3. This can be resolved in an extended Kalman filter where computed yaw and yaw rate are used as

output and input, respectively, in a straightforward manner.

- The sensors, in particular the gyro, have biases $\mathbf{b}$ and also scale errors $G$. The biases can partly be calibrated during stand still intervals but if the sensor is handheld the bias cannot be as easily identified.
- The platform has an unknown 3D orientation $\mathbf{q}$, so all IMMU measurements above have been multiplied with a rotation matrix $R(\mathbf{q})$ above. This would require an orientation filter, which is a standard problem from IMMU measurements Callmer (2011).
- The magnetic field is disturbed indoors, assumption 4, so the measured magnetic field at a point $p$ can be expressed as $\mathbf{y}_k^m = R(\mathbf{q})\left(\mathbf{m}^o + \mathbf{d}^m(p)\right)$, where $\mathbf{m}^o$ is the earth magnetic field and $\mathbf{d}^m(p)$ is a local position dependent disturbance. This can be mapped, see de Vries et al. (2009), but this requires a tremendous mapping effort, assumption 2.
- The measured acceleration is the sum of gravitation and linear acceleration $\mathbf{y}_k^a = R(\mathbf{q})\left(\mathbf{g} + \mathbf{a}_k\right)$ why it does not give an exact measure of Z. The alternatives are to either consider linear acceleration as a disturbance, or to use advanced gait models where the linear acceleration pattern in each step is learnt and compensated for.

At rest, the sensor orientation $R(\mathbf{q})$ can be estimated using $\mathbf{y}_k^a$, $\mathbf{y}_k^m$ and $\mathbf{z}_k^E$. We know that $\mathbf{y}_k^a$ is pointing up and that $\mathbf{z}_k^E$ is pointing east so by crossing $\mathbf{y}_k^a$ with $\mathbf{z}_k^E$ we get the vector $\mathbf{z}_k^N$ that is pointing north,

$$\mathbf{z}_k^N = \mathbf{y}_k^a \times \mathbf{z}_k^E. \tag{7}$$

Normalizing the vectors gives us the estimated rotation matrix

$$\begin{aligned} R(\hat{\mathbf{q}}) &= \left( \frac{\mathbf{z}_k^E}{\|\mathbf{z}_k^E\|} \quad \frac{\mathbf{z}_k^N}{\|\mathbf{z}_k^N\|} \quad \frac{\mathbf{y}_k^a}{\|\mathbf{y}_k^a\|} \right) \\ &= \left( \bar{\mathbf{z}}_k^E \quad \bar{\mathbf{z}}_k^N \quad \bar{\mathbf{y}}_k^a \right) \end{aligned} \tag{8}$$

In summary, the two fundamental yaw related observations are

$$\begin{aligned} y_k^{\dot{\psi}} &= (0,\ 0,\ 1) R(\mathbf{q})^T \mathbf{y}_k^\omega && (9)\\ &\approx (0,\ 0,\ 1) R(\hat{\mathbf{q}})^T \mathbf{y}_k^\omega && (10)\\ &= \bar{\mathbf{y}}_k^a \cdot \mathbf{y}_k^\omega && (11)\\ &= \bar{\mathbf{y}}_k^a \cdot \left( G^\omega \omega + \mathbf{b}^\omega \right) + e_k^{\dot{\psi}}, && (12)\\ \mathbf{z}_k^E &= \mathbf{y}_k^m \times \mathbf{y}_k^a && (13)\\ &= R(\mathbf{q}) \left( \left( \mathbf{m}^o + \mathbf{d}_k^m \right) \times \left( \mathbf{g} + \mathbf{a}_k \right) \right) + e_k^{\mathbf{z}}, && (14)\\ y_k^{\psi} &= \arctan(\mathbf{z}_k^E(2)/\mathbf{z}_k^E(1)) && (15)\\ &\approx \psi_k + e_k^{\psi} && (16) \end{aligned}$$

# 5 Adaptive Filtering

This section proposes an adaptive filter to estimate heading based on the two computed measurements $y_k^{\dot{\psi}}$ and $y_k^{\psi}$. However, it is good practice in filtering to also include mechanisms for divergence detection and restarting the filter with appropriate initial conditions. This will be devoted its own section.

A basic linear model suitable for a Kalman filter is based on the state $\psi_k$ and the model

$$\psi_{k+1} = \psi_k + T y_k^{\dot{\psi}} + e_k^{\dot{\psi}}, \tag{17}$$

$$y_k^{\psi} = \psi_k + e_k^{\psi} \tag{18}$$

where $T$ is gyro sampling time.

## 5.1 Gyro Sensor Error Modeling

To allow for trustworthy dead reckoning when $y_k^{\psi}$ is polluted by magnetic disturbances, the gyro sensor deficiencies described in Section 4 should be compensated for.

Introducing constant gyro error compensation terms $S^{\omega}$ and $\mathbf{c}^{\omega}$

$$\tilde{\mathbf{y}}_k^{\omega} = S^{\omega}\mathbf{y}_k^{\omega} + \mathbf{c}^{\omega} + \mathbf{e}_k^{\omega} \tag{19}$$

$$= S^{\omega}\left(G^{\omega}\omega + \mathbf{b}^{\omega}\right) + \mathbf{c}^{\omega} + \mathbf{e}_k^{\omega} \tag{20}$$

$$= S^{\omega}G^{\omega}\omega + (S^{\omega}\mathbf{b}^{\omega} + \mathbf{c}^{\omega}) + \mathbf{e}_k^{\omega} \tag{21}$$

where $e_k \sim \mathcal{N}(0, \sigma_{\omega}^2)$. If $S^{\omega}$ and $\mathbf{c}^{\omega}$ are accurately estimated, some of the gyro sensor deficiencies can be cancelled out.

We therefore add the error compensation terms to the state vector $x_k$

$$x_k = \begin{pmatrix} \psi_k & \bar{S}^{\omega} & \mathbf{c}^{\omega} \end{pmatrix}^T. \tag{22}$$

where $\bar{S}^{\omega}$ is a vector of the nonzero terms in the matrix $S^{\omega}$ that should be estimated.

The dynamic model of $\psi_k$ now becomes

$$\psi_{k+1} = \psi_k + T\bar{\mathbf{y}}_k^{a} \cdot \bar{\mathbf{y}}_k^{\omega} + e_k^{\dot{\psi}}. \tag{23}$$

## 5.2 Magnetic Disturbance Modeling

Now, magnetic disturbances come in bursts, and so do good data. This is verified by Figure 1. A suitable model for this is a discrete Hidden Markov Model (HMM) for the measurement noise, so

$$e_k^{\psi} \sim \mathcal{N}(0, \sigma(r_k)^2), \tag{24}$$

$$P(r_k = i | r_{k-1} = j) = \Pi_{ij}. \tag{25}$$

where $r_k$ is the mode at time $k$, in our case undisturbed (u) or disturbed (d). $\Pi_{ij}$ is the so called mode transition probability matrix indicating the probability of switching from mode $j$ to mode $i$.

### 5.3 Estimation System

The optimal filter under this assumption becomes a Kalman filter bank. It executes multiple filters in parallel, each with its own history $r^i_{1:k}$ of assumed disturbed or undisturbed data. For each new measurement, all $r^i_{1:k}$ are split in two: one assuming the new measurement is undisturbed and one assuming it is disturbed. This makes the number of mode history combinations $2^k$, and all should be taken into consideration.

The posterior distribution for $\hat{x}_k$ becomes

$$\hat{x}_k \in \sum_{i=1}^{2^k} w_k(i)\mathcal{N}\Big(\hat{x}_k(i), P_k(i)\Big) \tag{26}$$

where $\hat{x}_k(i)$ and $P_k(i)$ is the state estimate and covariance, respectively, for a Kalman filter executed using the sequence $r^i_{1:k}$. The relative weight $w_k(i)$ is computed from each individual Kalman filter as

$$-2\log(w_k(i)) \propto \sum_{l=1}^{k} \varepsilon_l(i)^T S_l^{-1}(i)\varepsilon_l(i). \tag{27}$$

where $\varepsilon_l(i)$ is the filter innovation $\varepsilon_l(i) = y^\psi_l - \hat{y}^\psi_l(i)$ and $S_l^{-1}(i)$ is the innovation covariance of mode sequence $i$ at time $l$. There are many feasible approximations proposed in literature to mitigate the exponential complexity, and Section 7 describes one.

## 6 Divergence Monitoring and Reinitialization

The mode estimation in the Kalman filter bank is based on the filter innovations meaning that a measurement $y^\psi_k$ that differs greatly from the anticipated measurement $\hat{y}^\psi_k$ will be deemed as disturbed and vice versa. If the filter has converged, this approach works very well.

On the other hand, if the filter has not converged correctly, $\hat{y}^\psi_k$ will not be close to $y^\psi_k$ even when $y^\psi_k$ is undisturbed. This causes undisturbed measurements to be deemed disturbed and only the occasional disturbed measurement that is close to $\hat{y}^\psi_k$ will be regarded as trustworthy. The filter has subsequently diverged and can no longer find its way back to the correct values of $x_k$.

To monitor for divergence, a parallel system, independent of the filter, can be used. It is based on a test statistic that has a known distribution when the measurements are undisturbed. If this test statistic indicates that the measurements

are undisturbed with a very high probability, we get a robust heading estimate. If this robust estimate deviates significantly from the optimal filter estimate, filter divergence can be concluded and the filter is reinitialized with the robust heading estimate.

We can construct two independent test statistics based on the measurements as

$$\lambda_{1,k} = y_k^{\psi} - y_{k-1}^{\psi} \tag{28}$$

$$\lambda_{2,k} = T y_{k-1}^{\dot{\psi}}. \tag{29}$$

Under the null hypothesis of undisturbed measurements, the test statistics have the distributions

$$\lambda_{1,k} \sim \mathcal{N}(\psi_k - \psi_{k-1}, 2\sigma_{\dot{\psi}}^2) \tag{30}$$

$$\lambda_{2,k} \sim \mathcal{N}(\psi_k - \psi_{k-1}, T^2\sigma_{\dot{\psi}}^2) \tag{31}$$

A new test statistic can be based on the difference between $\lambda_{1,k}$ and $\lambda_{2,k}$

$$\lambda_k = \lambda_{1,k} - \lambda_{2,k} \sim \mathcal{N}(0, 2\sigma_{\dot{\psi}}^2 + T^2\sigma_{\dot{\psi}}^2). \tag{32}$$

To get a stronger test, the test statistics can be monitored over a sliding window. Since $\lambda_k$ and $\lambda_{k-1}$ share a common measurement in $y_{k-1}^{\psi}$, they are not independent. To avoid such cross terms, we propose the following test statistic

$$\bar{\lambda}_k = \frac{1}{2\sigma_{\dot{\psi}}^2 + T^2\sigma_{\dot{\psi}}^2} \sum_{j=0}^{L-1} \lambda_{k-2j}^2 \sim \chi^2(L) \tag{33}$$

$\bar{\lambda}_k$ is $\chi^2(L)$ distributed, where $L$ is the length of the sliding window. A threshold $\alpha$ can be determined, so that if $\bar{\lambda}_k < \alpha$, $y_k^{\psi}$ is very likely undisturbed. The length of the window $L$, becomes a tradeoff between false positives and missed detections.

# 7 Implementation

So far, the resulting algorithm has been described on a conceptual level. This section will describe an implementation that is suitable for real-time processing on a feasible device. Several design choices and assumptions are proposed to reduce the execution time.

## 7.1 Models

In all of our field trials, the device was positioned more or less horizontally, reducing the need for orientation compensation. We present an algorithm below that is simplified accordingly. The accelerometer signal was lowpass filtered to reduce the impact of user movements.

Since only measurements in the plane were considered, the dimension of $x_k$ could be reduced. In the implementation the state vector is augmented with bias and

scale error terms for the yaw rate sensor

$$x_k = \begin{pmatrix} \psi_k & s & c \end{pmatrix}^T \tag{34}$$

making the dynamic model

$$\begin{aligned} x_{k+1} &= A_k x_k + G e_k^{\dot{\psi}} \\ \begin{pmatrix} \psi_{k+1} \\ s \\ c \end{pmatrix} &= \begin{pmatrix} 1 & T y_k^{\dot{\psi}} & T \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \psi_k \\ s \\ c \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} e_k^{\dot{\psi}} \end{aligned} \tag{35}$$

while the measurement model remains unchanged.

## 7.2 Filter Implementation, IMM

The full Kalman filter bank described in Section 5.3 has an exponentially growing complexity since all possible mode hypothesis trajectories have to be maintained. Since this is unmanageable in the long run, an implementation that should run in real time has to be simplified in some ways. One option is to prune the tree of mode history by removing the least probable branches, another is to remove the mode histories altogether.

In this work we chose to use the Interacting Multiple Model (IMM) filter, Bar-Shalom et al. (2001), in the implementation. The IMM does not maintain the full history of all mode estimates. Instead it only uses the latest estimate of the mode probabilities $\mu_k = (\mu_k^u,\ \mu_k^d)$ and merges the results of the undisturbed and disturbed filters.

To estimate $\mu_k$ the IMM filter uses $\mu_{k-1}$, the mode switching probability $\Pi_{ij}$ and the innovation of each mode $i$: $y_k^{\psi} - C\hat{x}_k^i$ with covariance $S_k^i$.

### Mode Switching Probability

The true mode switching probability $\Pi_{ij}$ can be estimated using a large amount of data. The resulting mode switching probability would be the same going from disturbed to undisturbed as from undisturbed to disturbed. In an implementation this is unsuitable. Falsely perceiving a disturbed measurement as undisturbed is a larger problem than the opposite why we state that going from disturbed to undisturbed is more unlikely than the inverse transition.

In the implementation, the mode switching probability was set as

$$\Pi = \begin{pmatrix} 0.9 & 0.1 \\ 0.005 & 0.995 \end{pmatrix},$$

i.e. the probability of switching from the undisturbed to disturbed mode is 10%, while the reverse mode transition has a probability of only 0.5%

In practice, $\Pi_{ij}$ has the effect of introducing some sort of dynamics in the system. If both modes were equally likely, every measurement with a low innovation would be perceived as undisturbed. The mode switching probability will instead

state that some transitions are very unusual, so before a mode switch is made, multiple undisturbed measurements might be required. This makes sense in our case since the magnetic disturbances are location dependent.

### Disturbance Distribution Model

The influence of $\mathbf{d}_k^m$ has to be modeled and since the IMM filter uses Kalman filter innovations, the disturbance has to be modeled as Gaussian. The disturbance mean is set to 0 since it can be both positive and negative, and since $\mathbf{d}_k^m$ can result in $y_k^\psi$ having any value $[-\pi\ \pi]$, $\sigma(d) \gg \sigma(u)$.

### Filter Equations

The algorithm of how to merge the estimates, the time update, the measurement update, the mode probability estimation and the filter reinitialization now follows for convenience.

When a new measurement $y_k^\psi$ has been acquired, the first step is to calculate the mixing probabilities $\{\mu_{k-1|k-1}^{ji}\}_{i,j=\{u,d\}}$

$$\mu_{k-1|k-1}^{ji} = \frac{\Pi_{ij}\mu_{k-1}^{j}}{\sum_{l=\{u,d\}}\Pi_{il}\mu_{k-1}^{l}}. \tag{36}$$

The mixed estimates $\{\hat{x}_{k-1|k-1}^{0i}\}_{i=\{u,d\}}$ and covariances $\{P_{k-1|k-1}^{0i}\}_{i=\{u,d\}}$ then become

$$\hat{x}_{k-1|k-1}^{0i} = \sum_{j=\{u,d\}} \mu_{k-1|k-1}^{ji}\hat{x}_{k-1|k-1}^{j} \tag{37}$$

$$\begin{aligned} P_{k-1|k-1}^{0i} = \sum_{j=\{u,d\}} \mu_{k-1|k-1}^{ji}[P_{k-1|k-1}^{j} + \\ + (\hat{x}_{k-1|k-1}^{j} - \hat{x}_{k-1|k-1}^{0j})(\hat{x}_{k-1|k-1}^{j} - \hat{x}_{k-1|k-1}^{0j})^T] \end{aligned} \tag{38}$$

These mixed estimates are now time updated using (35)

$$\begin{aligned} \hat{x}_{k|k-1}^{i} &= A_k\hat{x}_{k-1|k-1}^{0i} \\ P_{k|k-1}^{i} &= A_kP_{k-1|k-1}^{0i}A_k^T + GQG^T. \end{aligned} \tag{39}$$

The states are then updated with $y_k^\psi$ using standard Kalman filter measurement updates with the respective disturbed or undisturbed models.

$$\hat{x}_{k|k}^{i} = \hat{x}_{k|k-1}^{i} + K_k^i(y_k^\psi - C\hat{x}_{k|k-1}^{i}) \tag{40}$$

$$P_{k|k}^{i} = P_{k|k-1}^{i} - K_k^iS_k^iK_k^{iT} \tag{41}$$

$$S_k^i = CP_{k|k-1}^{i}C^T + R^i \tag{42}$$

$$K_k^i = P_{k|k-1}^{i}C^T(S_k^i)^{-1} \tag{43}$$

The mode probabilities $\mu_k^i$ are now updated as

$$\mu_k^i = \frac{\mathcal{N}(y_k^\psi; C\hat{x}_{k|k-1}^i, S_k^i) \sum_{j=\{u,d\}} \Pi_{ji} \mu_{k-1}^j}{\sum_{l=\{u,d\}} \mathcal{N}(y_k^\psi; C\hat{x}_{k|k-1}^l, S_k^l) \sum_{j=\{u,d\}} \Pi_{jl} \mu_{k-1}^j}. \tag{44}$$

The updated mode probabilities $\mu_k^i$ are finally used to calculate the overall estimate $\hat{x}_{k|k}$ and covariance $P_{k|k}$ as

$$\hat{x}_{k|k} = \mu_k^u \hat{x}_{k|k}^u + \mu_k^d \hat{x}_{k|k}^d \tag{45}$$

$$P_{k|k} = \sum_{i=\{u,d\}} \mu_k^i [P_{k|k}^i + (\hat{x}_{k|k}^i - \hat{x}_{k|k})(\hat{x}_{k|k}^i - \hat{x}_{k|k})^T]. \tag{46}$$

In the end, the filter estimates are used to detect filter divergence after $\bar{\lambda}_k$ has been updated with the latest measurements. So if $\mu_k^d > \beta$ while $\bar{\lambda}_k < \alpha$, the filter is deemed diverged and is reinitialized as

$$\hat{x}_{k|k} = \begin{pmatrix} y_k^\psi & 1 & 0 \end{pmatrix}. \tag{47}$$

To avoid having the IMM filter reinitialize due to the device being at rest for too long, a simple movement detector was used. If all dimensions of $\mathbf{y}_{k:k-N}^a$ and $\mathbf{y}_{k:k-N}^\omega$ were constant, the device was perceived as at rest and the measurements were ignored.

# 8 Experimental Results

The device that was used for data collection is a handheld IMMU in an HTC Sensation XE smartphone. The gyro sampling rate is chosen as 25 Hz, and the magnetometer sampling rate as 1 Hz. Our experimental results are based on 652 datasets. The 652 datasets constitute more than 20 hours of data and their combined walking trajectories cover roughly 80 km. The average length of each experiment is thus around 120 m, which in most cases provides only a small sample path of the complete building. That means that the variability of the tests is quite large, and a fraction of the tests has only good magnetometer data and another fraction has only bad data, just by coincidence. The test sites are selected as some 50 public buildings, such as shopping malls, spread over five continents, to get thorough experience of magnetic disturbances in practice.

We have used a dedicated smoothing filter based on both wifi and inertial measurements and a building map to obtain ground truth that is as accurate as possible. The resulting ground truth is very good, and the quality has been manually inspected for all 20 hours of data. The proposed heading estimator is evaluated in two ways: first, the assumptions of the high degree of magnetic disturbances is confirmed, and second, the performance of the estimation system is assessed.

To evaluate the results, the estimated yaw and the magnetic yaw measurements were compared to the ground truth, and the mean absolute error was computed

for each. Since the transients of the filter estimates should not be included in the comparison, the first third of each dataset was removed when the error was computed.

## 8.1 Detailed Evaluation

Two representative routes have been selected for detailed evaluation. One where the system performs well and one when it actually performs worse than the magnetic yaw measurements.

### Successful Experiment

The first route used for detailed experimental validation is the same as in Figure 1 but taken four laps. Each lap has four 90° turns and takes about 65 seconds to walk. A major magnetic disturbance is present each lap for about 45 seconds. So for a majority of the trajectory there is a severe magnetic disturbance. There is also a significant bias in the angular velocity measurements.

The results of the yaw estimation are shown in Figure 4. Despite severe magnetometer disturbances, the bias in the gyro measurement is accurately estimated and compensated for, Figure 6. The estimated yaw is robust against disturbances and quite accurate once the estimate has converged.

The mean absolute error of the magnetic yaw measurements is 1.08 rad (62°) of the last two thirds of the data, while the mean absolute error of the yaw estimate is 0.23 rad (13°).

For the filter to converge, we generally need two good data windows separated in time and heading. The covariance of the yaw and bias estimates are therefore significantly reduced at about 70 seconds, see Figures 4 and 6, which is when the second good data segment starts.

The mode probability estimation seems to be working quite well, Figure 5. The very disturbed sections are estimated to be disturbed with a very high probability once the filter has converged, see for example at around 100 seconds. The undisturbed sections also have a high probability of being estimated as undisturbed, shown at around 80 seconds and 150 seconds. Since the whole distribution of undisturbed data is overlapped by the distribution of the disturbed data, it is generally easier to classify the disturbed data.

The gain correction value is not properly excited by the system so the gain estimate covariance is not significantly reduced during the experiment. The final gain value is $s = 0.999$ which is very close to the initial value of 1.000.

### Unsuccessful Experiment

An unsuccessful experiment is shown i Figure 7. The magnetometer yaw measurements are often centered around the ground truth but they are very noisy. In the experiment, this results in the filter never accurately converging before the first disturbed segment around $t = 40$. This causes the yaw estimate to be

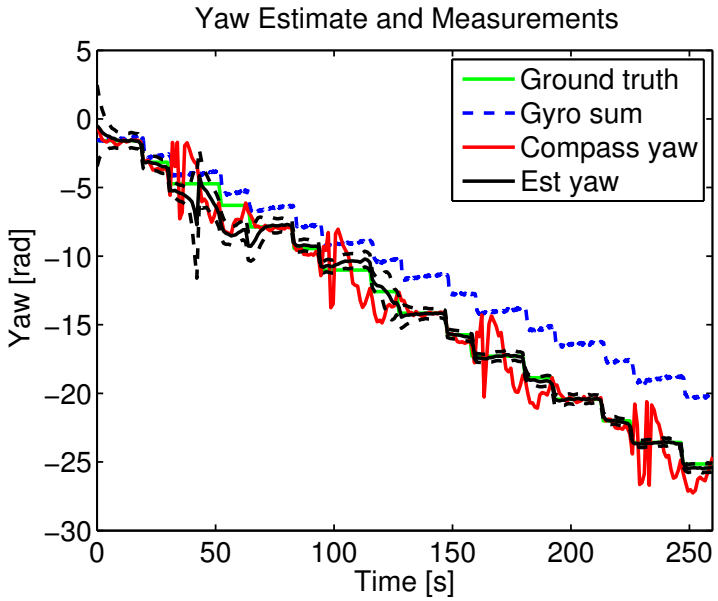


***Figure 4:*** *Magnetic yaw in solid red, ground truth in solid green, cumulative summation of gyro measurements in dashed blue and our estimate in solid black. Estimate covariance displayed using 2 $\sigma$ in dashed black. Note the periodicity of the magnetic disturbances.*

pulled away by bad magnetometer measurements and the bias is also incorrectly estimated. The filter has diverged and needs to be reinitiated.

Unfortunately, this data set is too short for getting a data window with good magnetometer data. Therefore, the backup system that monitors the difference between the gyro and the magnetometer measurements never perceives the magnetometer data as reliable, why no reinitialization is initiated.

When studied, the filter estimate is actually not that bad, since it is mainly suffering from an offset. Also the bias estimate has over time converged and in the end it eliminates the gyro bias quite well, Figure 9. In a positioning experiment the offset would of course result in the position estimate taking off in the slightly wrong direction. But, if an external positioning support of any kind is available, even if only temporarily, such an error could be detected and corrected.

The estimated mode probabilities, Figure 8 in general indicate that the magnetometer measurements are undisturbed even though they are quite uncertain. This uncertainty is often a good indication of an unlucky experiment.

As in the successful experiment, the gain estimate stays almost the same throughout the experiment and its covariance is not reduced.

The erroneous yaw estimate causes the mean absolute error to be large, 0.67 rad

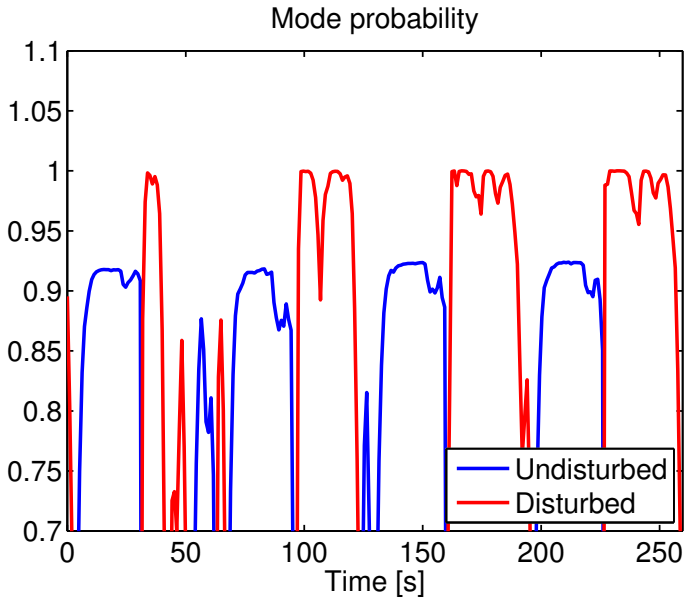


***Figure 5:*** *Estimated mode probabilities for well performing experiment.*

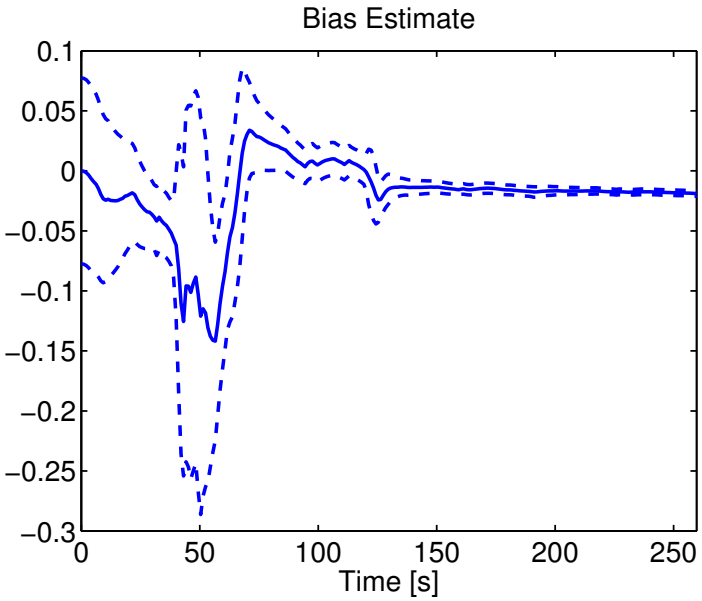


***Figure 6:*** *The estimated gyro bias with estimate covariance plotted with* $2\sigma$*. Final value is* $c = -0.0204$*.*

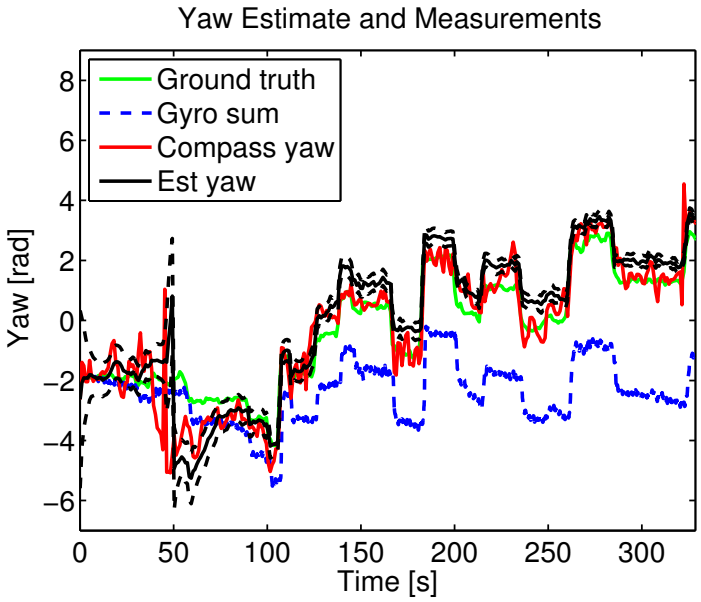


***Figure 7:*** *Magnetic yaw in solid red, ground truth in solid green, cumulative summation of gyro measurements in dashed blue and our estimate in solid black. Estimate covariance displayed using 2 $\sigma$ in dashed black. The estimation failed since the filter never managed to converge correctly.*

(39°). The magnetometer on the other hand does not have a large bias and is just very noisy most of the time. This results in the mean absolute error being 0.42 rad (24°) for the magnetic yaw for the same period.

## 8.2 Large Scale Evaluation

Using 652 datasets, different aspects of the heading estimation problem could be evaluated. We will first present statistics about the magnetic yaw measurements and then the performance of the proposed estimation system.

### Magnetic Disturbances

How often is the magnetic data actually disturbed? To assess this we studied how the magnetometer yaw measurements deviate from the ground truth. A yaw measurement is here deemed undisturbed if both the current and the previous measurement are good, in the sense that they deviate less than 0.35 rad (20°) from ground truth.

Using a total of more than 20 hours of data (more than 72000 samples from the magnetometer), our conclusion is that 46.4% of the magnetometer yaw measurements are disturbed. A histogram of all magnetometer yaw measurement errors is shown in Figure 10. Most measurements are centered around 0 and even though disturbances of more than 1.5 rad (86°) exist, they are quite rare. The

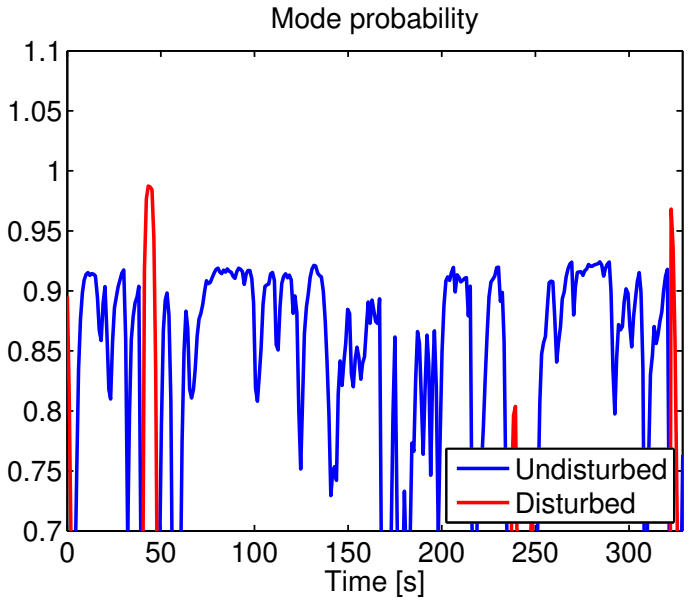


***Figure 8:*** *Estimated mode probabilities for unsuccessful experiment.*

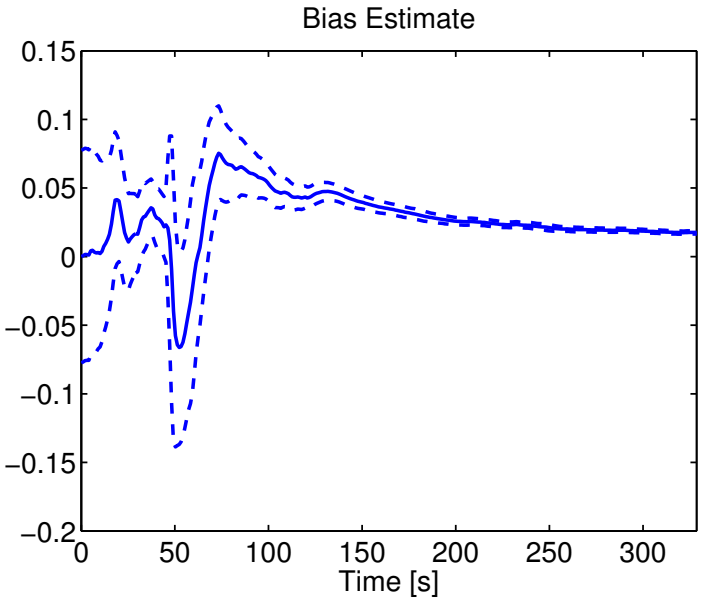


***Figure 9:*** *The estimated gyro bias with estimate covariance plotted with* $2\sigma$*. Final value is* $c = 0.0284$*.*

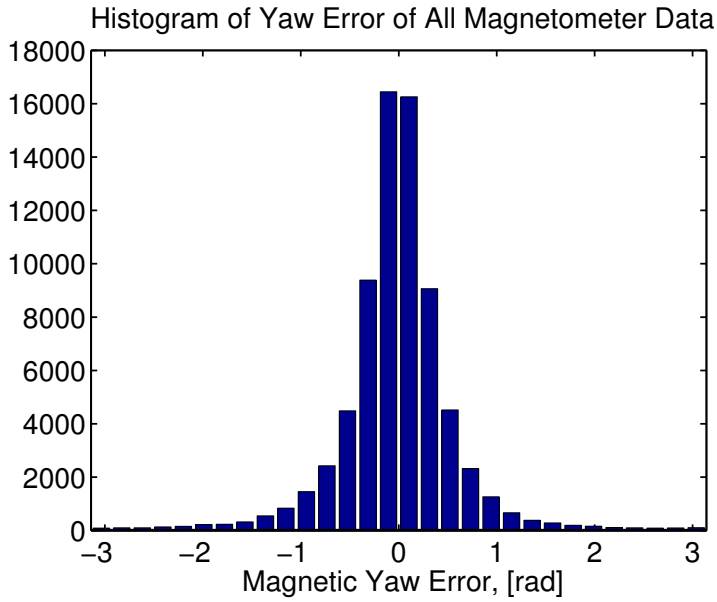


***Figure 10:*** *Histogram of the error in yaw of all magnetometer measurements.*

vast majority of all measurements are within ±1 rad (±57°) of the correct heading.

One could argue that modeling the disturbances in magnetic yaw should be done using a uniform distribution since it can be anything. If this was the case, the histogram in Figure 10 would have been the sum of a very narrow Gaussian distribution centered around 0 for all the correct measurements and a flat distribution on both sides of the Gaussian peak. Since this is not the case, the disturbances are actually better modeled using a Gaussian distribution with a very large covariance.

#### Yaw Estimation

The algorithm was evaluated using 652 datasets and the mean absolute error of the proposed algorithm is compared to the raw magnetometer signal for each dataset. A scatter plot of these means is shown in Figure 11. The result of each dataset is indicated by a black dot. The value on the $y$ axis shows the mean absolute error of the magnetometer yaw and the value on the $x$ axis shows the mean absolute error of the estimated yaw for each data set. Points above the red line shows where the estimate is better than the magnetometer (in the mean), and below this line the algorithm actually performs worse. In total, 548 experiments were successful while 104 were unsuccessful. 305 experiments reduced the error by at least 30% while only 23 increased the error by more than 30%.

In a few cases, the estimate is significantly worse due to very unfavorable magnetic data. The outlier dot in position [0.67 0.42] corresponds to the experiment

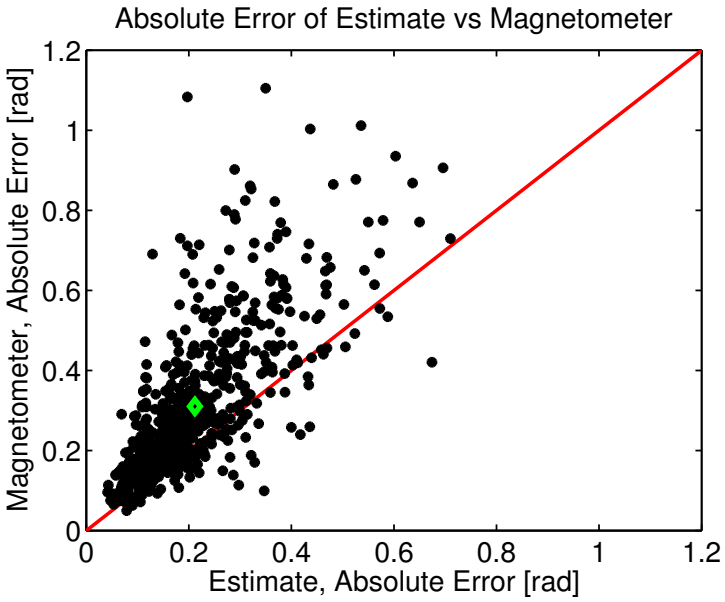


***Figure 11:*** *Scatter plot of mean absolute error of magnetometer yaw versus the mean absolute error of the filter estimates for 652 datasets. The green marker is the mean result of all datasets and all dots above the red line indicate a performance improvement and vice versa.*

that was shown in Section 8.1, which is not really bad but only biased. The reason for the fraction of seemingly poor estimates is that we choose to plot the mean error, which suffers from an estimation bias. The standard deviation of the error is in all cases much smaller for the estimate. Note that the bias is caused by the lack of data windows with good magnetometer measurements, which is a problem that would disappear for longer trajectories than our rather short ones.

The mean value, marked by a green diamond in Figure 11 indicates that the algorithm reduces the mean absolute error from 0.310 rad (17.8°) to 0.218 rad (12.5°) on average, a reduction of 29.8%.

#### Gyro Sensor Error Estimation

As indicated by the detailed experiments in Section 8.1, the gain correction term $s$ is not really excited by the system. In all datasets used in the large scale evaluation, the gain state covariance is more or less unchanged throughout the experiments and the final estimated value is always very close to the initial value 1.000. This suggests that the gain correction term is more or less insignificant why it can be removed.

On the other hand, the gyro bias term is most often accurately estimated. It significantly reduces the drift in estimated yaw during long segments of magnetic disturbances, as previously discussed in Section 8.1.

# 9 Conclusions

Common approaches to indoor positioning makes use of either radio fingerprinting or building information, or both. This requires that the venue must be known in advance, a case that does not apply to rescue personell. The dominating approach for positioning rescue personell is based on foot-mounted inertial magnetometer measurement units. The short moments of foot stand-still enables bias calibration of the IMMU, and dead reckoning gives only a linear drift in distance. However, heading estimation is a partly open problem, where a universal solution does not appear to exist.

We have proposed a framework based on an optimal filter that integrates gyroscope measurements and corrects this yaw value with computed compass direction from a magnetometer, where the noise is modelled as a discrete Hidden Markov Model. It is shown that the short moments where the magnetometer happens to be reliable is sufficient to learn the bias parameters in the gyroscope, and correct the drift.

We also propose a robust filter that compares the total yaw angle over a sliding window from the gyroscope and magnetometer, respectively. If they are consistent, a robust though not optimal, course estimate is obtained. This can then be compared to the optimal filter estimate for divergence detection and initialization after a filter restart.

We have verified our assumptions on disturbances with a couple of representative data sets, and the large scale performance based on a large number of field tests.

## Acknowledgments

This work was supported by CADICS, a Linnaeus center funded by the Swedish Research Council, and by the Excellence Center at Linköping-Lund in Information Technology (ELLIIT).

## Bibliography

E. Bachmann, X. Yun, and A. Brumfield. Investigation of the effects of magnetic variations on inertial/magnetic orientation sensors. *IEEE Robotics and Automation Magazine*, 14(3):76 – 87, 2007.

E. Bachmann, J. Calusdian, E. Hodgson, and X. Yun. In situ heading drift correction for human position tracking using foot-mounted inertial/magnetic sensors. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012.

Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. Wiley, New York, 2001.

J. Borenstein and L. Ojeda. Heuristic drift elimination for personnel tracking systems. *Journal of Navigation*, 63(3):591 – 606, 2010.

J. Borenstein, L. Ojeda, and S. Kwanmuang. Heuristic reduction of gyro drift in imu-based personnel tracking systems. In *Proceedings of SPIE, the International Society for Optical Engineering*, 2009.

J. Callmer. Topics in localization and mapping. Linköping University, Sweden, 2011. Licentiate Thesis No 1489.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors. *IEEE Transactions on Signal Processing*, 2013. Submitted.

W. H. K. de Vries, H. E. J. Veeger, C. T. M. Baten, and F. C. T. van der Helm. Magnetic distortion in motion labs, implications for validating inertial magnetic sensors. *Gait Posture*, 29(4):1786 – 1792, 2009.

E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

D. Gusenbauer, C. Isert, and J. Krösche. Self-contained indoor positioning on off-the-shelf mobile devices. In *Intl. Conf. on Indoor Positioning and Indoor Navigation IPIN*, 2010.

T. Harada, H. Uchino, T. Mori, and T. Sato. Portable absolute orientation estimation device with wireless network under acceleration situation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

C. W. Kang and C. G. Park. Improvement of attitude estimation using hidden markov model classification. In *AIAA Infotech@Aerospace 2010 Conference and Exhibit*, 2010.

J. K. Lee and E. J. Park. Minimum-order kalman filter with vector selector for accurate estimation of human body orientation. *IEEE Transactions on Robotics*, 25(5):1196 – 1201, 2009.

K. Pahlavan, F. Akgul, Y. Ye, T. Morgan, F. Alizadeh-Shabdiz, M. Heidari, and C. Steger. Taking positioning indoors – wi-fi localization and gnss. *InsideGNSS*, May 2010.

J. Rantakokko, J. Rydell, P. Strömbäck, P. Händel, J. Callmer, D. Törnqvist, F. Gustafsson, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: multisensor systems and cooperative localization. *Wireless Communications, IEEE*, 18(2):10–18, 2011.

D. Roetenberg, H. J Luinge, C. T. M. Baten, and P. H. Veltink. Compensation of magnetic disturbances improves inertial and magnetic sensing of human body segment orientation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3):395 – 405, 2005.

A. M. Sabatini. Quaternion-based extended kalman filter for determining orientation by inertial and magnetic sensing. *IEEE Transactions on Biomedical Engineering*, 53(7):1346 – 1356, 2006.

A. M. Sabatini. Estimating three-dimensional orientation of human body parts by inertial/magnetic sensing. *Sensors*, 11(2):1489 – 1525, 2011.

R. Stirling, J. Collin, K. Fyfe, and G. Lachapelle. An innovative shoe-mounted pedestrian navigation system. In *Proceedings of European Navigation Conference GNSS*, 2003.

Y. S. Suh, Y. S. Ro, and H. J. Kang. Quaternion-based indirect kalman filter discarding pitch and roll information contained in magnetic sensors. *IEEE Transactions on Instrumentation and Measurement*, 61(6):1786 – 1792, 2012.

X. Yun, E. Bachmann, and R. McGhee. A simplified quaternion-based algorithm for orientation estimation from earth gravity and magnetic field measurements. *IEEE Transactions on Instrumentation and Measurement*, 57(3): 638 – 650, 2008.

# Paper B

# Robust Heading Estimation Indoors using Convex Optimization

*Authors:* Jonas Callmer, David Törnqvist and Fredrik Gustafsson

*Edited version of the paper:*

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors using convex optimization. In *International Conference on Information Fusion*, 2013b. Submitted.

*Preliminary version:*

Technical Report LiTH-ISY-R-3061, Dept. of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden.

# Robust Heading Estimation Indoors using Convex Optimization

Jonas Callmer, David Törnqvist and Fredrik Gustafsson

Dept. of Electrical Engineering,
Linköping University,
SE–581 83 Linköping, Sweden
callmer@isy.liu.se

**Abstract**

The problem of estimating heading is central in the indoor positioning problem based on measurements from inertial measurement and magnetic units. Integrating rate of turn angular rate gives the heading with unknown initial condition and a linear drift over time, while the magnetometer gives absolute heading, but where long segments of data are useless in practice because of magnetic disturbances. A basic Kalman filter approach with outlier rejection has turned out to be difficult to use with high integrity. Here, we propose an approach based on convex optimization, where segments of good magnetometer data are separated from disturbed data and jointly fused with the yaw rate measurements. The optimization framework is flexible with many degrees of freedom in the modeling phase, and we outline one design. A recursive solution to the optimization is derived, which has a computational complexity comparable to the simplest possible Kalman filter. The performance is evaluated using data from a handheld smartphone for a large amount of indoor trajectories, and the result demonstrates that the method effectively resolves the magnetic disturbances.

## 1 Introduction

Dead reckoning systems are based on a movement and a direction in which the unit is moving. To determine the direction of a user, a magnetometer is often used. It works as a compass, measuring the earth magnetic field and thereby giving a measurement of the unit heading. Indoors, magnetic disturbances from steel structures, electrical wiring, cabinets etc, are often present making magnetometer readings unreliable.

A gyro measures rate of turn, which means that it measures how the heading is changing. It can be used to stabilize the heading estimation, but when used alone the integrated heading estimate is prone to drift off.

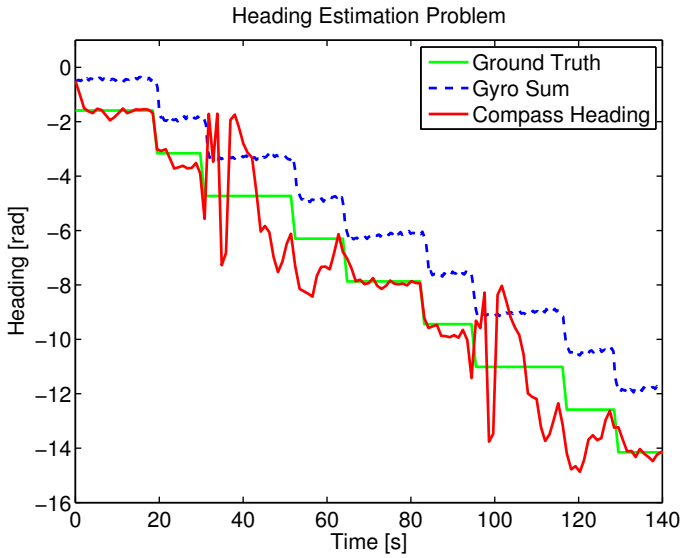


***Figure 1:*** *How do we fuse the disturbed magnetometer signal (solid red) with the biased gyro measurements (dashed blue) to reproduce the ground truth (solid light green)?*

In this work we study the problem of accurately estimating heading indoors using an inexact gyro and a magnetometer which is sometimes heavily disturbed.

The problem is illustrated in Figure 1. It shows the cumulative sum of gyro signals with the drift clearly present, the significantly disturbed magnetic heading and the actual ground truth heading. Note that sometimes the magnetic field is undisturbed, providing information of user heading that can and should be used. In this work we aim to support the heading estimation by detecting and utilizing these undisturbed magnetic heading sections.

The problem of how to handle magnetic disturbances for usage in positioning has not received a lot of attention. Often the signal has been discarded as unreliable indoors and simply ignored Stirling et al. (2003); Borenstein and Ojeda (2010).

The most common approach to handle the magnetic disturbances is to study the signal norm and/or the dip angle Sabatini (2006); Lee and Park (2009); Harada et al. (2004); Sabatini (2011). If either the signal norm or the dip angle deviates from a narrow interval, the magnetic signal is deemed completely unreliable and is subsequently discarded. In Roetenberg *et al* Roetenberg et al. (2005) this is handled a bit less abrupt by modeling the magnetometer measurement noise as a function of the signal norm and the dip angle.

Based on our large scale experiments in public venues, our conclusion is that the norm or the dip angle of the magnetic signal cannot be used to find good data segments in a robust way, or even to find just some of the bad ones. The

magnetometer signal simply fluctuates too much even during good data segments to be useful.

A slightly different approach was taken in Kang and Park (2010) where the magnetic disturbances are detected using a hidden Markov model where the modes represent disturbed and undisturbed data, respectively. The resulting mode estimate affects the parameter settings of a Kalman filter that estimated the attitude.

All in all, ways to detect and handle magnetic disturbances have not received a great deal of attention and the previously suggested solutions almost all focus on features in the magnetic signal itself.

Our suggested heading estimation system is based on the assumption that the errors in the gyro measurements are from insufficient calibration. The angular velocity signal is hence seen as primarily corrupted by sensor deficiencies such as a slightly incorrect gain and a small bias. By estimating bias and gain error correction terms and also an initial heading, the current heading can be estimated robustly.

The bias and gain correction parameters are estimated using convex optimization. They are estimated by altering the shape of the vector of the cumulative sum of angular velocity measurements, see the dashed blue signal in Figure 1. The aim is to tweak it until it matches the shape of the vector of magnetometer heading measurements as closely as possible. In the end, the problem is formulated as a weighted least squares problem.

A more standard approach to estimate heading would be to use a Kalman filter. The angular velocity measurements are used as an input in the time update while the magnetometer measurements are used in a measurement update. The disturbances could for example be handled by studying the innovation and discarding the perceived bad measurements.

There are two main shortcomings of such a Kalman filter solution:

- Disturbances are detected using only the latest magnetometer measurement.
- It needs to be reinitialized if the filter diverges.

The proposed method can handle both of these issues. Since it studies the whole trajectory of data, it becomes easier to detect and discard the disturbed sections of the magnetic signal. This makes the proposed method more robust. It will also recover automatically if diverged once more data is available, given that some of that data is undisturbed. Finally, we have derived a computationally cheap solution to the problem, that compares well with a Kalman filter.

If one still wants to use a Kalman filter to estimate heading, a more elaborate scheme than the standard Kalman filter is needed to handle the disturbances and to reinitiate if diverged.

In this work we have restricted the problem to yaw estimation in the 2D plane. The sensor is from a smartphone that is handheld and the method is evaluated

on a large number of data sets. Strengths and weaknesses of the method are also discussed.

## 2 Problem Fundamentals

The true heading $h(t)$ is a function of the true but unknown angular velocities $\omega(t)$ as

$$h(t) = \int_0^t \omega(\tau) d\tau + h(0). \tag{1}$$

Now we sample $\omega(t)$ using a perfect sensor with sampling time $T$ and zero order hold. That means $\omega_k = \omega(t_k)$. Assuming that the bandwidth of the user movements is small compared to the sampling frequency, the true time discrete heading $h_k$ is

$$h_k = T \sum_{i=1}^{N} \omega_i + h_0 \tag{2}$$

where $k = NT$.

Unfortunately, sensors are not perfect, and cheap MEMS gyros are prone to be poorly calibrated why their measurements have slight errors. In this work we have assumed that these errors are a minor constant change in gain $\bar{g}$, a small constant additive bias $\bar{b}$ and measurement noise $\bar{e}_k \sim \mathcal{N}(0, \bar{\sigma}_\omega^2)$ making the gyro measurements

$$y_k^\omega = \bar{g}\omega_k + \bar{b} + \bar{e}_k. \tag{3}$$

This can be rewritten as

$$\omega_k = g y_k^\omega + b + e_k \tag{4}$$

where $g = \frac{1}{\bar{g}}$ and $b = -\frac{\bar{b}}{\bar{g}}$ which are constant, and $e_k \sim \mathcal{N}(0, \sigma_\omega^2)$. Consequently, we have the fundamental assumption that

$$h_k = T \sum_{i=1}^{N} (g y_i^\omega + b + e_i) + h_0 \tag{5}$$

where the variance of $h_k$ is $NT^2\sigma_\omega^2$.

Now, $g$, $b$ and $h_0$ are estimated as $\hat{g}$, $\hat{b}$ and $\hat{h}_0$ and we assume that $\sigma_\omega$ is very small. The estimated heading is then

$$\begin{aligned} \hat{h}_k &= \sum_{i=1}^{N} T(\hat{g} y_i^\omega + \hat{b}) + \hat{h}_0 \\ &= \hat{g} \sum_{i=1}^{N} T y_i^\omega + \hat{b} \sum_{i=1}^{N} T + \hat{h}_0. \end{aligned} \tag{6}$$

Introduce $m_k = \sum_{i=1}^{N} T y_i^\omega$ and $n_k = \sum_{i=1}^{N} T$. Since they are given by the sensor and therefore are known, $m_k = m_{k-1} + T y_k^\omega$ and $n_k = n_{k-1} + T$. The heading estimation model is now

$$\hat{h}_k = m_k \hat{g} + n_k \hat{b} + \hat{h}_0. \tag{7}$$

Putting the sought after states in a vector $\theta$

$$\theta = \begin{pmatrix} g & b & h_0 \end{pmatrix}^T, \tag{8}$$

the estimated heading at time $k$ is

$$\begin{aligned} \hat{h}_k &= \begin{pmatrix} m_k & n_k & 1 \end{pmatrix} \hat{\theta} \\ &= \Omega_k \hat{\theta}. \end{aligned} \tag{9}$$

Since $\theta$ is assumed constant, the vector of all heading estimates $\hat{H}_{0:k} = \begin{pmatrix} \hat{h}_0 & \hat{h}_1 & \dots & \hat{h}_k \end{pmatrix}^T$ can be written as

$$\begin{aligned} \hat{H}_{0:k} &= \begin{pmatrix} 0 & 0 & 1 \\ m_1 & n_1 & 1 \\ & \vdots & \\ m_k & n_k & 1 \end{pmatrix} \hat{\theta} \\ &= \Omega_{0:k} \hat{\theta} \end{aligned} \tag{10}$$

# 3 Heading Estimation

To estimate $\theta$ we use all the heading measurements acquired by the magnetometer from time 0 to $k$, $\bar{Y}_{0:k}^h = \begin{pmatrix} y_0^h & y_1^h & \dots & y_k^h \end{pmatrix}^T$. Some of these will be greatly disturbed while some will be correct.

The approach is based on optimization and estimates $\theta$ by altering the shape of $\hat{H}_{0:k}$. The sought after shape is the one that matches $\hat{H}_{0:k}$ to the seemingly accurate sections of $\bar{Y}_{0:k}^h$ as closely as possible. To focus the attention only to those sections of the data, the shape matching is performed using weighted least squares.

## 3.1 Detect Disturbed Magnetometer Readings

We therefore need to detect which parts of the magnetometer data that are disturbed and which ones that are not. One measure of the trustworthiness of the magnetometer data is to compare how well the magnetometer and the gyro measurements agree. The residual

$$\lambda_k = y_k^h - y_{k-1}^h - T y_{k-1}^\omega \tag{11}$$

is only close to 0 if the magnetometer and the gyro measure the same change in heading over that time interval. The interval is only one sample here but can

be generalized to more samples by summing up measurements if the sampling times differ.

Sometimes disturbed magnetometer data and gyro data coincide over a short interval just by chance, making $\lambda_k$ small when we do not want it to. This will make disturbed measurements seem trustworthy. To avoid the problems this can introduce and also to capture the fact that disturbances are location dependent and therefore time dependent, $\lambda_k$ is averaged using a sliding window.

$$\bar{\lambda}_k = \frac{1}{N} \sum_{i=k-N+1}^{k} |\lambda_i| \tag{12}$$

The length of the window $N$ is a tradeoff between having false positives and false negatives. We averaged the absolute value of $\lambda_k$ to avoid problems when $\lambda_k$ was changing sign in the interval. Since $\lambda_0$ is unknown it is given a value reflecting a poor measurement.

The actual weight function used in the vector shape matching is

$$w(\bar{\lambda}_k) = e^{-\bar{\lambda}_k^2/s}. \tag{13}$$

It is close to 1 when $\bar{\lambda}_k$ is small and to 0 when $\bar{\lambda}_k$ is large. $s$ determines the shape of $w(\bar{\lambda}_k)$, thereby stating what a "small" or a "large" value of $\bar{\lambda}_k$ is.

In the end, the weight function is quite ad hoc. It makes sense when studied step by step, but we cannot in any way prove that it is the best choice of weight. It is reasonable that the weights should be large when the magnetometer and gyro seem to agree. It is also reasonable that we should smooth out the measured difference between the sensors a bit over time, since disturbances are often location dependent. $\bar{\lambda}_k$ should then be transformed so that if $\bar{\lambda}_k$ is large $w(\bar{\lambda}_k) \sim 0$ and if $\bar{\lambda}_k$ is small $w(\bar{\lambda}_k) \sim 1$. The actual values used on $N$ and $s$ have been determined by studying a number of challenging datasets. At the end of the day, what is important is getting a weight vector that looks reasonable.

The weight function is included in the optimization problem as a diagonal matrix $W(\bar{\lambda}_{0:k})$ with $\{w(\bar{\lambda}_i)\}_{i=0}^{k}$ as its entries.

## 3.2 Parameter Estimation

The parameter estimation can be solved as a weighted least squares problem. We know that the approximate true value of the parameters are $\hat{g} \approx 1$ and $\hat{b} \approx 0$, so we should put a restriction on their deviation from these values. Since there are no hard boundaries on these values, giving fixed intervals for the values as constraints is not a good idea.

The parameter values should instead be kept in order using regularization. That means that the deviation of the parameter value from the assumed one is associated with an additional cost in the cost function. It is not obvious though how to add regularization terms to the cost function, since the dimension of $\hat{H}_{0:k}$ is constantly growing and therefore also the cost is growing.

There are two basic ways to solve this problem: normalizing the regularization terms using the actual cost or using the dimension of $\hat{H}_{0:k}$. In the first we solve the optimization problem with no restrictions on $\hat{b}$ and $\hat{g}$ as

$$C_k = \text{minimize}_{\hat{b}_k,\hat{g}_k,\hat{h}_{0,k}} \|\bar{Y}^h_{0:k} - \hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})\|^2_{W(\bar{\lambda}_{0:k})} \tag{14}$$

and note the final cost $C_k$. We then resolve the problem with the additional regularization terms that depend on $C_k$ and the parameter values.

$$\begin{aligned}\text{minimize}_{\hat{b}_k,\hat{g}_k,\hat{h}_{0,k}} &\|\bar{Y}^h_{0:t} - \hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})\|^2_{W(\bar{\lambda}_{0:k})} \\ &+ f_g^{C_k}(C_k, \hat{g}_k) + f_b^{C_k}(C_k, \hat{b}_k)\end{aligned} \tag{15}$$

An alternative would be to normalize the regularization terms using the dimension of $\bar{Y}_{0:k}$. Unfortunately, how the cost $C_k$ grows with the dimension will be different for each dataset, why such a normalization will be very approximative.

In this work we therefore normalized using the cost $C_k$ and used the regularization terms

$$f_g^{C_k}(C_k, \hat{g}_k) = \beta C_k(\hat{g}_k - 1)^2 \tag{16}$$

$$f_b^{C_k}(C_k, \hat{b}_k) = \alpha C_k \hat{b}_k^2 \tag{17}$$

Both have a quadratic cost for deviating from the nominal values of $\hat{b}_k$ and $\hat{g}_k$. This means that we allow small changes of $\hat{b}_k$ and $\hat{g}_k$, while large changes are associated with great costs. Also, the values of $\alpha$ and $\beta$ are chosen so that these costs relate well to each other. Since the deviation in $\hat{b}_k$ is likely to be larger than in $\hat{g}_k$, we use $\beta \gg \alpha$. Their actual values have in the end been chosen by studying the experimental results of a couple of datasets.

## 3.3 Magnetic Heading Vector Unwrapping

Before one can solve (15) there is a crucial preprocessing step. In order to match $\bar{Y}^h_{0:k}$ and $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$ they have to look like the gyro sum and the magnetic heading in Figure 1. When measured, the raw magnetometer heading measurements $Y^h_{0:k} \in [-\pi\ \pi]$ but such a vector cannot be straightforwardly compared to $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$. First it has to be unwrapped.

Unfortunately, it is extremely important that the unwrapping is done correctly. For the estimation to work, it is crucial that the magnetic heading is unwrapped so that it is centered around the ground truth at all times, like in Figure 1. If it is not, $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$ will be tweaked to match a vector that has the wrong shape, ruining the estimation.

Of course, we do not have the ground truth to do the unwrapping around. Instead we have to use the gyro sum for this. We add multiples of $2\pi$ to the magnetic vector as $\bar{Y}^h_i = Y^h_i + j \cdot 2\pi$, $j = \{..., -1, 0, 1, ...\}$, so that $-\pi \leq \bar{Y}^h_i - \hat{H}_i(0, 1, Y^h_0) \leq \pi$ for all $i$. This works fine if the dataset is short and/or the gyro bias is small. Else, eventually the gyro drift will have caused $\hat{H}_i(0, 1, Y^h_0)$ to drift off, resulting in

an incorrectly unwrapped magnetic heading vector and therefore poor heading estimation.

The alternative of unwrapping around $\hat{H}_i(0, 1, Y_0^h)$ is to unwrap using $\hat{H}_i(\hat{b}_{k-1}, \hat{g}_{k-1}, \hat{h}_{0,k-1})$. This works perfectly fine as long as $\hat{\theta}_{k-1}$ is estimated correctly. But, if there is an error in the estimation this will result in $Y_{0:k}^h$ being incorrectly unwrapped. When later used in estimation, the incorrectly unwrapped $\bar{Y}_{0:k}^h$ will make $\hat{\theta}_k$ even more wrong. The system hence becomes unstable. Primarily this can happen in datasets where large magnetic disturbances occur early.

Since unwrapping using $\hat{H}_i(0, 1, Y_0^h)$ works fine for short datasets and using $\hat{H}_i(\hat{b}_{k-1}, \hat{g}_{k-1}, \hat{h}_{0,k-1})$ works well for long datasets, we use them like that. For the first 30 seconds of data, the raw gyro sum measurements are used for unwrapping. When more than 30 seconds of data is available, the estimated heading vector is used. This has turned out to be a stable way to unwrap the magnetic heading vector as later experiments will show.

## 3.4 Solver Outline

Solving (15) once $\bar{Y}_{0:k}^h$ is correctly unwrapped is straightforward since it is just a weighted least squares problem. Even though the dimensions of $\bar{Y}_{0:k}^h$, $\hat{H}_{0:k}$ and $W(\bar{\lambda}_{0:k})$ grow with time, the solving time is constant using Cholesky decomposition.

Differentiating the cost function and putting it equal to zero results in (14) being solved as

$$A_k \hat{\theta}_k = b_k \tag{18}$$

where $A_k$ is a symmetric 3x3 matrix and $b_k$ is 3x1. Using Cholesky decomposition, $A_k = L_k L_k^T$, (18) can be solved very cheaply. Over time, $A_k$ also evolves in a structured way as $A_k = A_{k-1} + x_k x_k^T$ where

$$x_k = \begin{pmatrix} m_k w_k^{1/2} & n_k w^{1/2} & w_k^{1/2} \end{pmatrix}^T \tag{19}$$

where $w_k = w(\bar{\lambda}_k)$. Therefore, the Cholesky decomposition only has to be done once and then it is updated with each new measurement using rank-one updates, $L_k L_k^T = L_{k-1} L_{k-1}^T + x_k x_k^T$

Calculating the actual cost $C_k$ is also cheap since the terms in the cost function evolve in a structured way.

$$\begin{aligned} C_k &= (\bar{Y}_{0:k} - \Omega_{0:k}\hat{\theta}_k)^T W_{0:k} (\bar{Y}_{0:k} - \Omega_{0:k}\hat{\theta}_k) \\ &= \bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k} - 2\hat{\theta}_k^T \Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k} + \hat{\theta}_k^T \Omega_{0:k}^T W_{0:k} \Omega_{0:k} \hat{\theta}_k \\ &= \underbrace{\bar{Y}_{0:k-1}^T W_{0:k-1} \bar{Y}_{0:k-1} + \bar{y}_k^2 w_k}_{\bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k}} \end{aligned}$$

$$- 2\hat{\theta}_k^T \underbrace{\left(\Omega_{0:k-1}^T W_{0:k-1} \bar{Y}_{0:k-1} + \bar{y}_k w_k \Omega_k^T\right)}_{\Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k}}$$
$$+ \hat{\theta}_k^T \underbrace{\left(\Omega_{0:k-1} W_{0:k-1} \Omega_{0:k-1} + x_{k+1} x_{k+1}^T\right)}_{\Omega_{0:k}^T W_{0:k} \Omega_{0:k}} \hat{\theta}_k \tag{20}$$

where $W_{0:k} = W(\bar{\lambda}_{0:k})$ and $w_k = w(\bar{\lambda}_k)$. So if $\bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k}$, $\Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k}$ and $\Omega_{0:k}^T W_{0:k} \Omega_{0:k}$ are saved for the next iteration, updating the cost function will then only require multiplication of 3-dimensional vectors and 3x3 dimensional matrices.

Now, when (15) is differentiated one gets

$$(A_k + \bar{C}_k)\hat{\theta}_k = \bar{b}_k \tag{21}$$

where

$$\bar{C}_k = \begin{pmatrix} 2C_k\beta & 0 & 0 \\ 0 & 2C_k\alpha & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{22}$$

$A_k + \bar{C}_k$ can be subsequently be Cholesky decomposed as $L_k L_k^T + c_{\alpha,k} c_{\alpha,k}^T + c_{\beta,k} c_{\beta,k}^T$ using two more rank-one updates where

$$c_{\beta,k} = \begin{pmatrix} \sqrt{2\beta C_k} & 0 & 0 \end{pmatrix}^T$$
$$c_{\alpha,k} = \begin{pmatrix} 0 & \sqrt{2\alpha C_k} & 0 \end{pmatrix}^T \tag{23}$$

Thus, (15) can be solved very cheaply no matter the dimension of $\bar{Y}_{0:k}^h$, $\hat{H}_{0:k}$ and $W(\bar{\lambda}_{0:k})$.

# 4 Experimental Results

The devices used for data collection are inertial magnetic measurement units from hand held smartphones. Sensors used in smartphones are not so well calibrated, since calibration is an expensive procedure. Therefore, gyro bias and scale errors are common problems in these low grade sensors. We will present the detailed results from one experiment and then mass evaluations using a large number of datasets.

When a new magnetometer heading measurement $Y_k$ is available, $\hat{\theta}_k$ is re-estimated using all data. $\hat{h}_k$ is then computed as (9) and stored in a vector $\bar{H}_{0:k} = [\bar{H}_{0:k-1} \ \hat{h}_k]$. $\bar{H}_{0:k}$ is referred to as the filtering estimate since it for each time instant contains the estimate that was given at that time.

Later, when more data is available, improved estimates of earlier headings $\hat{H}_{0:k-1}$ are readily available from (10), since $\theta$ is assumed constant throughout the ex-

periment. $\hat{H}_{0:\bar{k}}(\hat{\theta}_{\bar{k}})$ where $\bar{k}$ is the final time instant, will be referred to as the smoothing estimate.

In almost all datasets, $\hat{H}_{0:\bar{k}}(\hat{\theta}_{\bar{k}})$ will be a better estimate than $\bar{H}_{0:\bar{k}}$. In a few rare cases, the final magnetometer measurements pull the estimate away from the true values. This causes the smoothing trajectory to be off during a larger part of the trajectory, making it slightly worse than the filtering version.

## 4.1 Detailed Experiment

In the detailed dataset an HTC Sensation XE smartphone was used for data collection. A gyro sampling rate of 25 Hz and a magnetometer sampling rate of 1 Hz were used. In (11), this difference in sampling time has been handled by summing up all the angular velocity measurements between the magnetometer measurements.

The route used for experimental validation is the same as in Figure 1. It has four turns of 90 degrees and each lap takes about 65 seconds to walk. A major magnetic disturbance is present each lap for about 45 seconds. That means, a majority of the trajectory is constituted of a severe magnetic disturbance. There is also a significant bias in the angular velocity measurements.

The filtering results are shown in solid black in Figure 2. In the top plot it is clear that two segments of good data, one in the beginning and one starting around $t = 70$ seconds, are needed to stabilize the results. Before the second segment the parameter estimates, bottom plot, fluctuate quite a lot, but after around 70 seconds they become more stable. The magnetometer disturbances starting around $t = 95$ therefore have a much smaller influence on the filter estimates than the disturbances starting around $t = 30$.

Also included are the smoothing results in dashed magenta. The main difference between these and the filtering results is that also the early disturbances can now be handled making the smoothing estimate close to ground truth for the entire experiment. The final estimated parameter values are $\hat{g} = 1.005$, $\hat{b} = -0.020$ and $\hat{h}_0 = -1.668$, Figure 4. The mean absolute error is 0.924 rad for the magnetic heading, 0.476 rad for the filtering result and 0.243 for the smoothing result.

Figure 3 shows the weights calculated as in (13). It is clear that the weights are the largest and most consistent for the good magnetometer data sections around $t = 20$, $t = 80$ and $t = 140$. During the disturbed periods the weights are most often low. There are some spikes in the weights in sections where the magnetometer disturbances happen to look like the gyro data, like at $t = 120$, but this is more or less unavoidable. Their influence on the result is small in the end.

## 4.2 Mass Experiments

We have also conducted mass experiments, 651 datasets from more than 50 venues in 5 continents, to evaluate the performance of the heading estimation system. For each dataset, the estimated heading is compared to the ground truth that has

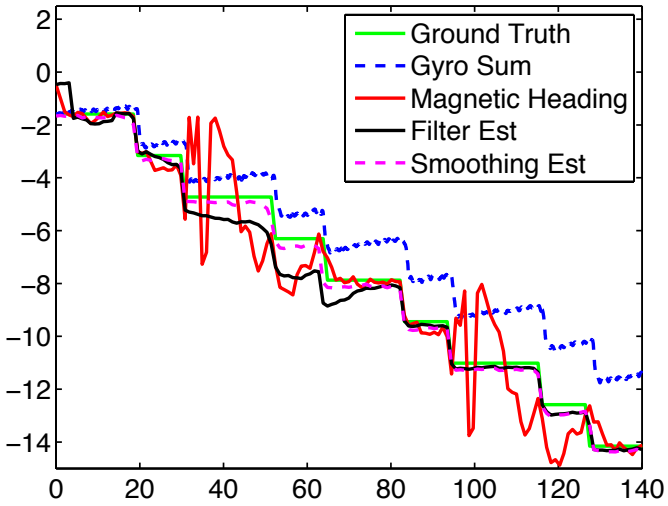


***Figure 2:*** *Magnetic heading in solid red, ground truth in solid light green, cumulative summation of gyro measurements in dashed blue, the filter heading estimate in solid black and the smoothing heading estimate in dashed magenta. The filter estimate becomes stable after 75 seconds while the smoothing estimate is stable for the entire data set.*

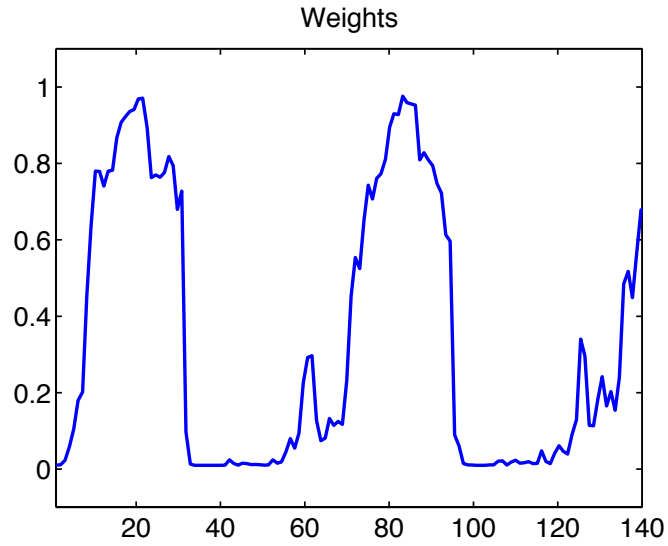


***Figure 3:*** *The weights assigned to the heading measurements used for heading estimation. For the most time the good magnetometer heading sections have a high weight and the disturbed ones have a low.*

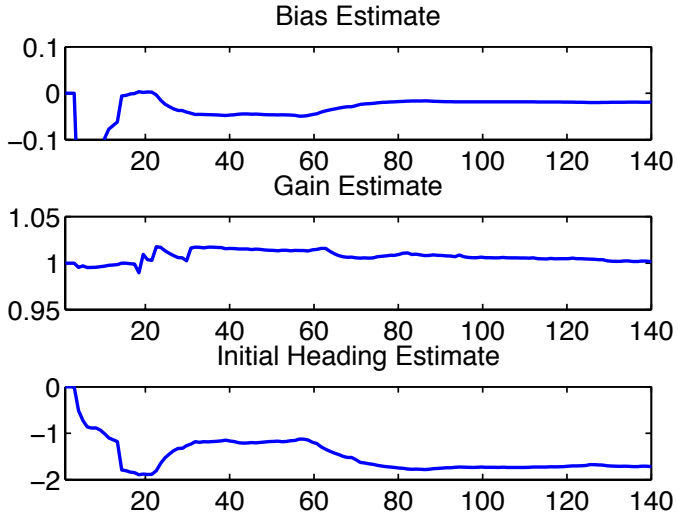


***Figure 4:*** *The estimated bias, gain and initial heading parameters. After the second good magnetometer data segment is found around* $t = 80$*, the parameters converge to their final values:* $\hat{b} = -0.020$*,* $\hat{g} = 1.005$ *and* $\hat{h}_0 = -1.668$*.*

been created based on knowledge about how the user was walking. The datasets are between 15 seconds and 8 minutes long.

The results are shown as mean absolute error for each dataset. In Figure 5, each experiment is indicated as a dot. The value on the $x$ axis is the mean absolute error of the filter result while the value on the $y$ axis is the mean absolute error of the magnetic heading. The red line indicate whether an improvement has occurred or not. All dots above the line are experiments where the estimate is better than the magnetic heading, while all dots below the line mean that the result has worsened. The green diamond is the mean result which is 0.334 rad for the magnetic heading, 0.266 for the filtering and 0.208 for the smoothing.

Figure 5 show the filtering results. That means, for each time instant only information up until that time instant is used to estimate the heading. In Figure 6, the smoothing estimate is shown. That means that all data has been used to estimate all the headings. The smoothing results are better since there are fewer dots under the red line, and for most of the experiments the dots are moved to the left compared to the filtering results, indicating a reduction in mean absolute estimation error.

In practice, the main difference between the filtering results and the smoothing results is in the beginning of the dataset. After the estimation has converged, the filtering and smoothing results are more or less identical.

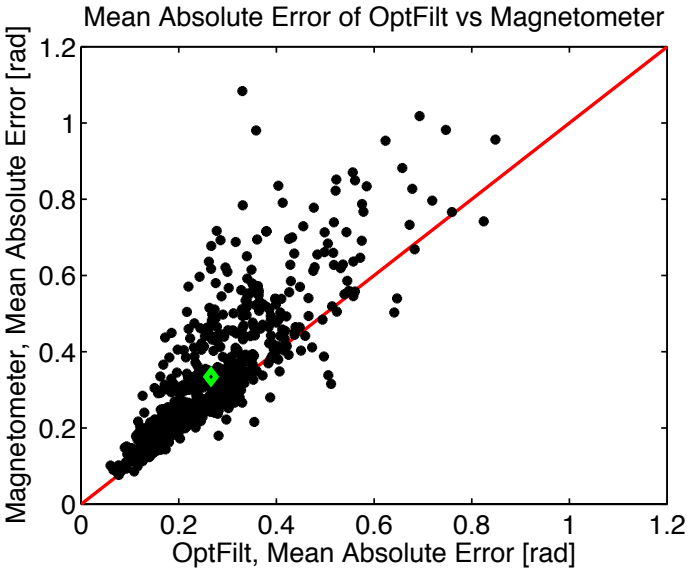


***Figure 5:*** *Mean absolute error of the filtering result plotted versus the mean absolute error of the magnetic heading. Each dot represent one dataset. All dots above the line indicate the result has improved and the further to the left the smaller the error.*

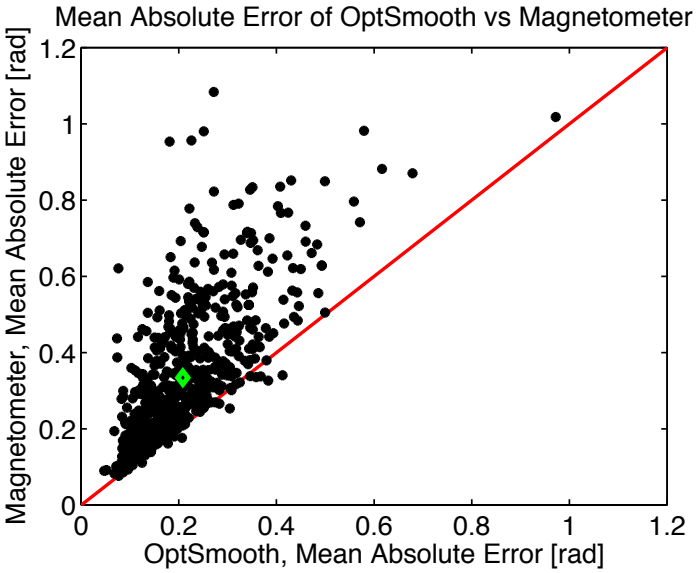


***Figure 6:*** *Mean absolute error of the smoothing result and the magnetic heading. Compared to Figure 5 most of the dots have moved to the left, since the smoothing results are better than the filtering results.*

The experiments show that the system can produce reliable estimates of heading even under heavy magnetic disturbances and that it can recover from poor early estimates.

# 5 Discussion

We have presented a method to estimate heading indoors. It is based on the assumption that the gyro measurements are correct down to a slight multiplicative error and an additive error. These errors are estimated from the magnetic heading measurements using weighted least squares. The experimental results indicate that the method can produce robust accurate estimates of user heading. We will now discuss the pros and cons of the system and how easy it is to work with.

The method uses the entire data set to reestimate the parameters at each iteration. An advantage with this approach is that earlier mistakes can be corrected once more data is available. Computationally this method is also extremely cheap and the number of computations for each iteration is constant throughout the dataset, Section 3.4.

Convex optimization can be a powerful tool in solving large scale signal processing problems due to the impressive solvers available, but sometimes fitting the problem into the solver framework requires a bit of ad hoc tampering. This led to, in our case, that the weighted optimization solution using batched data was not that easy to trim. Some terms like the weights and the regularization terms are chosen quite ad hoc.

In the end, the solver is only looking for the cheapest solution given the cost function and it can sometimes be quite hard to figure out why an occasional poor solution can be so cheap and what can be done to adjust the cost function to avoid it. Even though all steps seem reasonable when they are made, a system where they are all put together can turn out to be hard to trim. Therefore care has to be taken when designing a convex optimization based solution to a signal processing problem, but in the end the rewards can be great.

## 5.1 Conclusions

The fundamental assumption that the summation of the angular velocities is correct down to a few parameters, seems viable. The errors in estimated heading acquired in the experiments are commonly small and the system can handle significant magnetic disturbances.

Convex optimization offers powerful solutions to signal processing problems but one has to be careful when designing the problem formulation. The method is somewhat less forgiving than for example the Kalman filter but its ability to, among others, handle large sets of data compensates for that kind of limitations.

## Acknowledgments

The authors would like to thank prof Stephen Boyd at Stanford University and Daniel Petersson at Automatic Control, Linköping University for helpful discussions.

This work was supported by CADICS, a Linnaeus center funded by the Swedish Research Council, and by the Excellence Center at Linköping-Lund in Information Technology (ELLIIT).

## Bibliography

J. Borenstein and L. Ojeda. Heuristic drift elimination for personnel tracking systems. *Journal of Navigation*, 63(3):591 – 606, 2010.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors using convex optimization. In *International Conference on Information Fusion*, 2013. Submitted.

T. Harada, H. Uchino, T. Mori, and T. Sato. Portable absolute orientation estimation device with wireless network under acceleration situation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

C. W. Kang and C. G. Park. Improvement of attitude estimation using hidden markov model classification. In *AIAA Infotech@Aerospace 2010 Conference and Exhibit*, 2010.

J. K. Lee and E. J. Park. Minimum-order kalman filter with vector selector for accurate estimation of human body orientation. *IEEE Transactions on Robotics*, 25(5):1196 – 1201, 2009.

D. Roetenberg, H. J Luinge, C. T. M. Baten, and P. H. Veltink. Compensation of magnetic disturbances improves inertial and magnetic sensing of human body segment orientation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3):395 – 405, 2005.

A. M. Sabatini. Quaternion-based extended kalman filter for determining orientation by inertial and magnetic sensing. *IEEE Transactions on Biomedical Engineering*, 53(7):1346 – 1356, 2006.

A. M. Sabatini. Estimating three-dimensional orientation of human body parts by inertial/magnetic sensing. *Sensors*, 11(2):1489 – 1525, 2011.

R. Stirling, J. Collin, K. Fyfe, and G. Lachapelle. An innovative shoe-mounted pedestrian navigation system. In *Proceedings of European Navigation Conference GNSS*, 2003.

# Paper C

# An Inertial Navigation Framework for Indoor Positioning with Robust Heading

*Authors:* Jonas Callmer, David Törnqvist and Fredrik Gustafsson

*Edited version of the paper:*

J. Callmer, D. Törnqvist, and F. Gustafsson. An inertial navigation framework for indoor positioning with robust heading. *IEEE Transactions on Instrumentation and Measurement*, 2013c. Submitted.

# An Inertial Navigation Framework for Indoor Positioning with Robust Heading

Jonas Callmer, David Törnqvist and Fredrik Gustafsson

Dept. of Electrical Engineering,
Linköping University,
SE–581 83 Linköping, Sweden
callmer@isy.liu.se

**Abstract**

Indoor localization in unknown environments is considered, using inertial measurements from accelerometers, gyroscopes and magnetometers. Foot-mounted inertial sensors allow for stand-still detection triggering zero velocity updates that reduces the inertial navigation system (INS) drift in distance traveled from cubical to linear in time. We present a statistical framework, based on an navigation model. The standard stand-still mode is complemented with binary modes of magnetic disturbances. Test statistics for these two mode estimation problems are derived. Instead of making hard decisions, a hidden Markov model filter is used to compute the mode probabilities, leading to soft measurement updates in the Kalman filter.

Based on this, a robust smoothed heading estimate is computed in a second stage using the magnetometer. The final position estimate is then obtained by fusing the INS output with the robust heading in a standard dead-reckoning filter. Experiments demonstrate that the robust heading decreases the relative error in position from 10% to less than 1%, despite large magnetic disturbances.

## 1 Introduction

Indoor positioning is an enabling technology for many applications. There is one branch of research based on knowledge of the infrastructure, such as Wifi networks Woodman and Harle (2009); Seitz et al. (2010); Li et al. (2006); Honkavirta et al. (2009); Bahl and Padmanabhan (2000) or building maps Woodman and Harle (2008); Widyawan et al. (2008); Walder et al. (2009); Krach and Robertson (2008). For professional users such as firefighters, police officers or soldiers, an indoor positioning system should be completely autonomous and not require any *a priori* knowledge of the building. This application is thoroughly described in Rantakokko et al. (2011), where many related references are found.

A foot mounted inertial magnetic measurement unit (IMMU) dead reckoning system is a common suggestion in literature, Foxlin (2005); Stirling et al. (2003); Ojeda and Borenstein (2007); Godha et al. (2006); Beauregard (2007). The IMMU contains a three axis accelerometer, a three axis gyroscope and a three axis magnetometer. The gyro measurements are integrated to keep track of sensor orientation and the accelerometer measurements are double integrated to give position changes. Position computed in this way suffers from a cubic drift in time, due to gravity leakage from small orientation errors. If one can detect every time the foot is solidly on the ground, a so called zero velocity update can be made Foxlin (2005); Stirling et al. (2003); Ojeda and Borenstein (2007); Godha et al. (2006); Beauregard (2007); Skog et al. (2010). The accumulated velocity errors are thereby corrected about once a second during walking, why the error growth in distance travelled can be reduced to linear in time. The total distance can be quite accurately determined, but the accumulated heading error is more severe and needs special treatment.

To remove the drift in heading a magnetometer can be used. The earth magnetic field gives the direction to the magnetic north pole, hence the heading. Unfortunately, magnetic fields are often disturbed indoors Bachmann et al. (2007); de Vries et al. (2009) due to steel structures and wiring. In orientation estimation, magnetic disturbances are commonly handled by studying the magnetic signal norm or the dip angle Harada et al. (2004); Lee and Park (2009); Sabatini (2011); Kang and Park (2010); Roetenberg et al. (2005). In Callmer et al. (2013c) it is showed that this approach has its limitations when applied to indoor positioning.

In this work we suggest a three step algorithm for incorporating the magnetic heading measurements in a foot mounted INS. First, the user trajectory is estimated using accelerometer, gyro and stand still updates only, similar to Foxlin (2005); Stirling et al. (2003); Ojeda and Borenstein (2007); Godha et al. (2006); Beauregard (2007). Second, a novel real time robust heading smoother is derived that estimates the heading using all magnetic heading measurements of the experiment. Third, the user positions are recomputed using the estimated heading vector.

Magnetic heading has been used to reduce drift in heading before, but then the experiments where either performed outdoors Yun et al. (2012) or the magnetic disturbances were not handled at all Jiménez et al. (2010). Sometimes, magnetic heading has just been deemed too unreliable for utilization indoors Stirling et al. (2003).

To remove the drift in heading for a foot mounted INS, *ad hoc* approaches have been suggested. Extensive calibration using pre-walked paths Bachmann et al. (2012); Bebek et al. (2010) has been shown to reduce the error significantly. This is though more of a postponement than a solution to the problem. Assumptions on user behavior have been made, like assuming the user is walking in straight lines Borenstein et al. (2009); Jiménez et al. (2010), or is primarily walking in four to eight pre-specified building dependent directions Borenstein and Ojeda

(2010); Abdulrahim et al. (2011). Assuming the user is walking the same paths multiple times Angermann and Robertson (2012), or having the user place out beacons along the way Renaudin et al. (2007) can also be used to reduce the built up errors.

We show that magnetic heading can be incorporated successfully into an indoor dead reckoning system, despite severe deterministic magnetic disturbances present. Experiments show that the drift in heading can be removed, significantly reducing the positioning error.

The paper is outlined as follows. Section 2 presents an inertial navigation system. Hidden Markov models can be used to provide additional information to the INS and such a framework is presented in Section 3. An example of such additional information is a novel approach to stand still detection with soft decisions, Section 4. Experimental positioning results using such an inertial navigation+stand still correction system are presented. Section 5 proposes an original real-time smoothing estimate of heading robust to magnetic disturbances and how to incorporate it into the inertial navigation+stand still system. Experimental positioning results of the full system are presented in Section 6 followed by conclusions in Section 7.

# 2 An Inertial Navigation Framework

We will now give an introduction to an inertial navigation system (INS) using an IMMU. The system estimates user velocity and position by integrating the acceleration and angular velocity measurements from the IMMU. To reduce the accumulated errors, the IMMU is mounted on the foot.

## 2.1 Principles of IMMU Based Dead Reckoning

The accelerometer measurements $\mathbf{y}_{a,k}$ are the specific force experienced by the sensor, which is the sum of the user accelerations and the normal force countering the gravity component. To estimate the user movements, the gravity component must be removed from the acceleration measurements. In order to do this correctly, the exact sensor orientation must be known. Unfortunately, the orientation is unknown and must be estimated. The orientation estimate is maintained using a gyro measuring the angular velocities, $\mathbf{y}_{\omega,k}$.

A slight error in orientation results in gravity component leaking into the perceived user acceleration. Since the gravity component is often large compared to user accelerations, leakage will quickly cause large errors in estimated user velocity and position.

To reduce the positioning errors of such a system, the velocity estimation is anchored at each step using zero velocity updates. By detecting that the foot is on the ground, a virtual measurement of zero velocity can be fed to the estimation system. This corrects the built up errors in velocity, and reduces the errors in position and orientation significantly. More on this in Section 4.

## 2.2 States and Inputs

The inertial navigation model is based on ten dynamical states. Position and velocity are estimated in all three dimensions, $\mathbf{p}_k$ and $\mathbf{v}_k$ respectively, and to accurately integrate the user accelerations, the sensor orientation is also estimated. The orientation is represented using unit quaternions, $\mathbf{q}_k$, which is a four dimensional vector. For vectors in a Cartesian coordinate system, a superindex is sometimes used to separate the elements in vectors, so $\mathbf{p} = (\mathbf{p}^x \ \mathbf{p}^y \ \mathbf{p}^z)^T$. The state vector is

$$\mathbf{x}_k = \begin{pmatrix} \mathbf{p}_k^T & \mathbf{v}_k^T & \mathbf{q}_k^T \end{pmatrix}^T. \tag{1}$$

It is quite common to also include bias states, particularly for the gyro Foxlin (2005); Abdulrahim et al. (2011); Jiménez et al. (2010). As indicated in Nilsson et al. (2012), the system errors experienced are most often not on the sensor level but on the modeling level. Among others they detected a drift in heading that is sensor bias independent but depend on which foot the sensor is on. We have therefore chosen not to include any sensor bias terms in the state vector.

The accelerometer $\mathbf{y}_{a,k}$ and gyroscope $\mathbf{y}_{\omega,k}$ measurements are used as direct inputs $u_k$ in the model,

$$\mathbf{u}_k = \begin{pmatrix} \mathbf{y}_{a,k}^T & \mathbf{y}_{\omega,k}^T \end{pmatrix}^T. \tag{2}$$

## 2.3 Dynamic Model

A discrete time dynamic model based on a first order Taylor expansion of the rotation dynamics and an assumption that the input is piecewise constant between the time instants $t_k = kT$ is given by

$$\begin{aligned} \begin{pmatrix} \mathbf{p}_{k+1} \\ \mathbf{v}_{k+1} \\ \mathbf{q}_{k+1} \end{pmatrix} &= \begin{pmatrix} I & TI & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \begin{pmatrix} \mathbf{p}_k \\ \mathbf{v}_k \\ \mathbf{q}_k \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2} I \\ TI \\ 0 \end{pmatrix} \Big( R^T(\mathbf{q}_k) \big( \mathbf{y}_{a,k} + \mathbf{r}_{a,k} \big) - \mathbf{g} \Big) \\ &\quad + \frac{T}{2} \begin{pmatrix} 0 \\ 0 \\ S'(\mathbf{q}_k) \end{pmatrix} \big( \mathbf{y}_{\omega,k} + \mathbf{r}_{\omega,k} \big). \end{aligned} \tag{3}$$

Here, $I$ denotes a $3 \times 3$ identity matrix and $\mathbf{g}$ is the gravity vector $\mathbf{g} = (0 \ 0 \ 9.82)^T$. The matrix $R^T(\mathbf{q}_k)$ rotates the local measurements of acceleration to global acceleration measurements using the quaternion estimates, and $S'(\mathbf{q}_k)$ converts the gyro measurements into changes in orientation, see Kuipers (1999); Callmer (2011).

The model error terms $\mathbf{r}_a \sim \mathcal{N}(\mathbf{0}, \sigma_a^2 I)$ and $\mathbf{r}_\omega \sim \mathcal{N}(\mathbf{0}, \sigma_\omega^2 I)$ denote the measurement noises from the accelerometer and gyro, respectively. Of these, $\mathbf{r}_\omega$ and $\mathbf{r}_a^z$ are the most important ones when it comes to affecting the filter estimates. $\mathbf{r}_\omega$ translates into uncertainty about the sensor orientation which means possible leakage from gravity component to user accelerations. If $\mathbf{r}_\omega$ is large, covariance of position estimate becomes very large.

Since the gravity component is subtracted, it is important that the measured gravity component is correct to avoid a drift in elevation. Experiments indicate though that the measured $\mathbf{g}$ differs between the static and the dynamic case. Therefore, $\mathbf{g}$ cannot be measured at stand still and later used in the dynamic model. $\mathbf{r}_a^z$ should therefore incorporate the uncertainty in the measured value of $\mathbf{g}$, to ensure a smooth estimated user trajectory.

The discrete time model can be iterated on its own, which corresponds to a simulation of a dynamical system, or repeated time updates in a Kalman filter (KF) framework. Since the discrete time model is based on some approximations, the sample interval $T$ should be chosen small. Usually, inertial measurements come in the order of hundreds of Hz, which is sufficient to neglect the approximations. Of course, additional information is needed, which will be added as Kalman filter measurement updates.

# 3 Exogenous Information Framework

Integrating inertial measurements give a large drift in orientation and therefore position, why additional information is needed to stabilize the drift.

## 3.1 Information Sources

We will use two kinds of exogenous information:

- Stand still detection, which triggers a measurement update based on zero velocity and zero specific force (no acceleration of the device). The latter is translated to that the accelerometer only measures gravity.
- Earth magnetic field disturbance detection. As will be demonstrated, indoor environments have a high degree of magnetic field disturbances caused by iron structures and electrical currents. However, occasionally, the magnetometer senses the earth magnetic field, and this can be used to correct heading if such magnetic measurements are correctly detected. In particular, true north direction of the earth magnetic field will be detected.

These information sources will be embedded in a common hidden Markov model framework described below.

## 3.2 Discrete Hidden Markov Model Framework

We will use the same framework for both detection problems, based on a discrete hidden Markov model (HMM) with two states representing useful and discardable information. The transition probability matrices are denoted $\Pi^S$ and $\Pi^M$, respectively, where for instance $\Pi^S_{00}$ is the probability that the stand still detection system remains in stand still mode. $S$ and $M$ will denote stand still and magnetic, respectively, for the remainder of this work. Figure 1 illustrates the common model.

The discrete mode state is denoted $\delta^S$ and $\delta^M$, respectively. The 'normal' mode

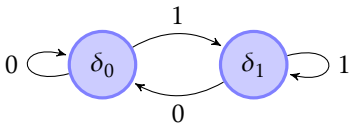


***Figure 1:*** *The two mode HMMs used in the stand still detection system and the magnetic disturbance detection system. $\Pi_{ii}$ is the probability of staying in mode $i$ and $\Pi_{ij}$ is the probability of transition from mode $i$ to mode $j$.*

is $\delta = 0$, which denotes the null hypothesis. However, we will not advocate a detection approach, since hard decisions in a Kalman filter framework may easily lead to divergence.

## 3.3 Test Statistics

We will in the sequel derive two test statistics with known distributions, and based on these propose observations of the latent discrete states. The test statistics is denoted $\lambda_k^S = \lambda^S(y_k)$ for the stand still mode, and analogously for the earth magnetic field mode. The test statistic has under the null hypothesis an analytical distribution

$$p(\lambda|H_0) = p(\lambda|\delta = 0) = p_0(\lambda). \tag{4}$$

The alternate hypothesis does not have an analytical form, as is usual in detection problems. Instead, we have empirically estimated the distributions based on large amounts of experimental data from different indoor venues, to obtain

$$p(\lambda|H_1) = p(\lambda|\delta = 1) = \hat{p}_1(\lambda). \tag{5}$$

The observation model for the HMM process can thus be stated as

$$p(\lambda|\delta) = \frac{(1-\delta)p_0(\lambda) + \delta\hat{p}_1(\lambda)}{\hat{p}_1(\lambda) + p_0(\lambda)} \tag{6}$$

## 3.4 Optimal HMM Filter

Since the distributions of the test statistics are known, each subproblem can be solved with an optimal filter. At each time $k$, the posterior mode probabilities $\mu_k^i = p(\delta_k = i|y_{1:k})$ of the discrete state $\delta_k$ is available. This corresponds to soft decisions about the two competing hypotheses, and no hard decisions need to be taken in this framework.

The mode probabilities $\mu_k^i$ at time $k$ are calculated recursively as

$$\begin{aligned}\mu_k^i &= P(\delta_k = i|y_k) \propto p(y_k|\delta_k = i)P(\delta_k = i|y_{k-1}) \\ &= p(\lambda_k|\delta_k = i)\sum_{j=1}^{N_\delta} \Pi_{ji}\mu_{k-1}^j.\end{aligned} \tag{7}$$

where $\delta_k$ is the mode state and $N_\delta$ are the number of modes, in this case $N_\delta = 2$.

Hence we have

$$\mu_k^i = \frac{p(\lambda_k|\delta_k = i)\sum_{j=1}^{N_r} \Pi_{ji}\mu_{k-1}^j}{\sum_{l=1}^{N_\delta} p(\lambda_k|\delta_k = l)\sum_{j=1}^{N_\delta} \Pi_{jl}\mu_{k-1}^j}. \tag{8}$$

The HMM filter is summarized in Algorithm 4.

**Algorithm 4** HMM filter

**Require:** Distributions $p_0(\lambda_k)$ and $p_1(\lambda_k)$ of test statistic $\lambda_k$. Initial mode probability $\mu_0$. Mode transition probability matrix $\Pi_{ij}$.

1: **for** $k = 1, \ldots, k_{\text{end}}$ **do**
2: Compute test statistic $\lambda_k$.
3: Update unnormalized likelihoods,

$$\begin{aligned} P_{0,k} &= p_0(\lambda_k)\left(\Pi_{00}\mu_{k-1}^0 + \Pi_{01}\mu_{k-1}^1\right) \\ P_{1,k} &= p_1(\lambda_k)\left(\Pi_{10}\mu_{k-1}^0 + \Pi_{11}\mu_{k-1}^1\right) \end{aligned} \tag{9}$$

4: Compute mode probabilities,

$$\mu_k^0 = \frac{P_{0,k}}{P_{0,k} + P_{1,k}} \qquad \mu_k^1 = \frac{P_{1,k}}{P_{0,k} + P_{1,k}}$$

5: **end for**

# 4 Stand Still Detection

The stand still detection system estimates the probability of stand still using the acceleration and angular velocity measurements registered by the IMMU and an HMM. Part of the stand still detection system has previously been presented in Callmer (2011); Callmer et al. (2010); Rantakokko et al. (2011).

## 4.1 Test Statistics

The test statistic is based on the accelerometer and gyro signal magnitudes to make it rotation invariant.

$$\lambda_k^S = \frac{\|\mathbf{y}_{a,k}\|^2}{\sigma_a^2} + \frac{\|\mathbf{y}_{\omega,k}\|^2}{\sigma_\omega^2} \tag{10}$$

where $\lambda_k^S \sim \chi^2(6, \beta)$ during stand still, $\beta = g^2/\sigma_a^2$. It has a non-central chi-square distribution since $\mathbf{y}_{a,k}$ has nonzero mean when the foot is stationary and six degrees of freedom since $\mathbf{y}_{a,k}$ and $\mathbf{y}_{\omega,k}$ are both three dimensional.

The movement distribution has been approximated empirically as a Gaussian mixture using a histogram of large amounts of experimental data.

## 4.2 Mode Switch Probability

We have used the mode transition probability matrix

$$\Pi^S = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix} \tag{11}$$

which states that the probability of going from stand still to moving or vice versa, is 2%. During walking each foot takes about one step per second resulting in roughly two mode transitions per second. So for a sensor sampling at 100 Hz, 2% of measurements will be a mode switch.

## 4.3 Stand Still Measurement Models

When a stand still has been detected, certain features are assumed true and used as measurement. First of all, the velocity is zero in all three dimensions,

$$\mathbf{0} = \mathbf{v}_k + \mathbf{r}_{v,k}. \tag{12}$$

In this work the zero velocity measurement noise is depending on the mode probability, $\mathbf{r}_{v,k} \sim \mathcal{N}(\mathbf{0}, \Sigma_v(\mu_k^S))$. If $\mu_k^S$ is large, $\Sigma_v(\mu_k^S)$ becomes small and vice versa.

Also, at a stand still, $\mathbf{y}_{a,k}$ should only contain the gravity component. Thereby it gives a measurement of the sensor orientation as

$$\mathbf{y}_{a,k} = R(\mathbf{q}_k)\mathbf{g} + \mathbf{r}_{g,k}. \tag{13}$$

where the noise model is analogous to above.

## 4.4 Stand Still Measurement Update

Each detected stand still should trigger a zero velocity update to ensure that the position estimates are not drifting while the user is standing still.

For other measurements, one and only one measurement update should be triggered for each stand still period. This should preferably be at the moment when the stand still probability is the highest. For reasons that will be clear soon, for the best stand still moments we introduce a second index $\bar{k}$ that is the index of each step. Subsequently, $\bar{k}$ is event based.

In Section 5 a method for heading estimation is presented that is based on the changes in gyro based heading and magnetic heading. This change is studied from stance to stance why we only want one orientation estimate from each stand still phase. To get the best possible estimate of orientation from each stance, the state estimates after the stand still update with the largest $\mu_k^s$, are chosen.

Other measurements like the gravitational measurement discussed in 4.3 should be used in the same way. It seems to contain a slight bias at each stand still that is individual for each stance. Using (13) multiple times during a stand still phase could therefore result in a slight bias in the orientation estimate.

## 4.5 Experiments

We will illustrate the basic framework with two complementary benchmark experiments. These were taken using a foot mounted semi-grade IMMU: a MicroStrain 3DM-GX3-25. The signals were sampled at 100 Hz. The states were estimated using and extended Kalman filter.

#### Experiment 1

In this experiment, the user has walked around in the corridors in Building B at Campus Valla, Linköping university. The user has also switched floor a number of times. The 2D positioning results are presented in Figure 2a.

The true start and end position were the same and basically all straight parts of the trajectory should be aligned with the $X$ and $Y$ axes. The final error in position is just over 40 meters for a trajectory of 540 meters, or 8%. The cause of this large error in position estimate is a slight bias in yaw and that some 90 degree turns are over- or underestimated . Since the heading drifted off, so did the position estimate.

#### Experiment 2

The second experiment, Figure 2b, is one where the user has been walking through a corridor, visiting every office first along one side and then every office on the other side on the way back. The actual start and end positions were the same and the straight paths in the beginning and in the end should be the same. The resulting error in final position estimate was about 25 meters. Since the experiment was 270 meters long, the error is about 10%.

# 5 Robust Heading for Indoor Positioning

Over time the gyro based heading drifts off resulting in the position estimates drifting off, Figure 2. To solve it, the heading estimation needs support from additional sensors.

Figure 3 shows the experimental heading estimates $\hat{\psi}_{0:\bar{k}}$ of the positioning system, plotted with the magnetic heading measurements $y^{\psi}_{0:\bar{k}}$. Each estimate and measurement is from an individual stand still phase, from the moment with the largest stand still probability $\mu^{s}_{k}$.

Three things are clear:

- The magnetometer signal is sometimes very noisy and disturbed.
- The shape of the gyro based heading estimate vector is similar to the shape of the magnetic heading vector.
- The gyro only heading estimate drifts over time.

In this case, Experiment 1 has a lot of magnetic disturbances, Figure 3a. Experiment 2, the one in the offices, is almost free from severe disturbances, Figure 3b.

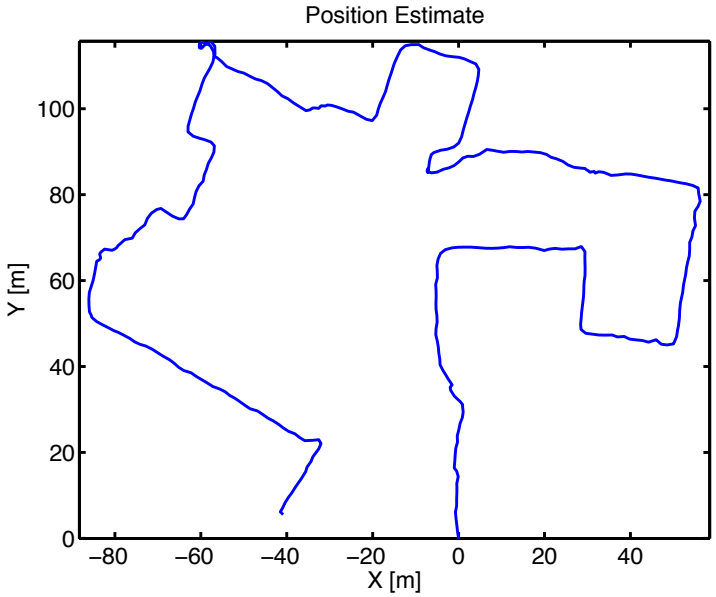


***(a)*** *Experiment 1. The user has walked along a series of corridors where most are perpendicular to each other.*

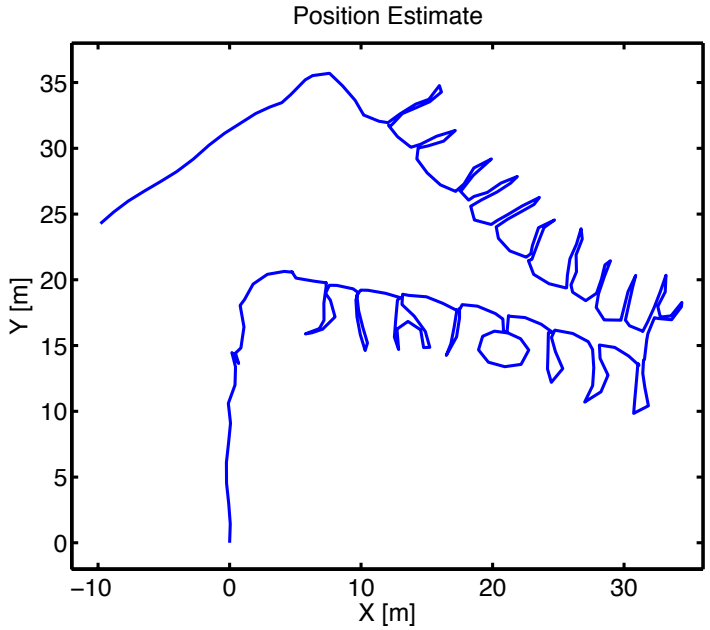


***(b)*** *Experiment 2. The user has walked in a corridor back and forth, visiting a lot of offices along the way.*

***Figure 2:*** *Experimental results of localization using only accelerometer and gyro data. A slight drift in heading translates to a large error in position since start and end position are the same in both experiments.*

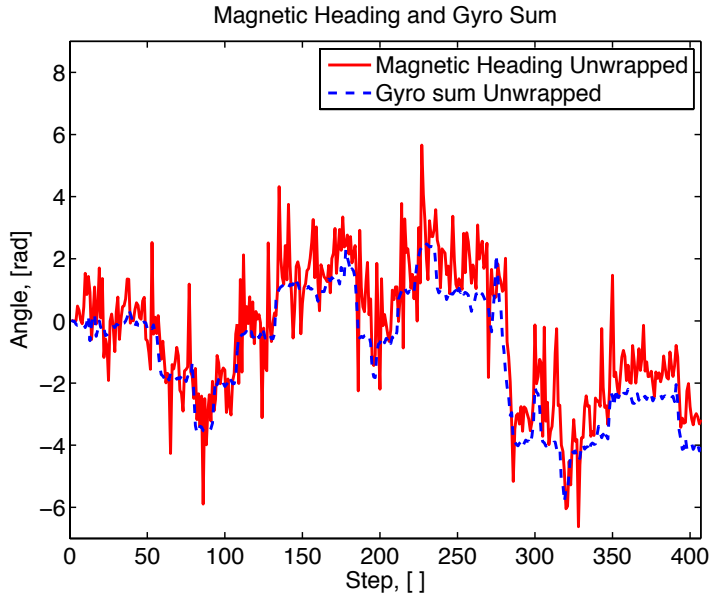


*(a) Experiment 1. Note how common and large the magnetic disturbances are. Final error in gyro heading is a little less than one radian.*

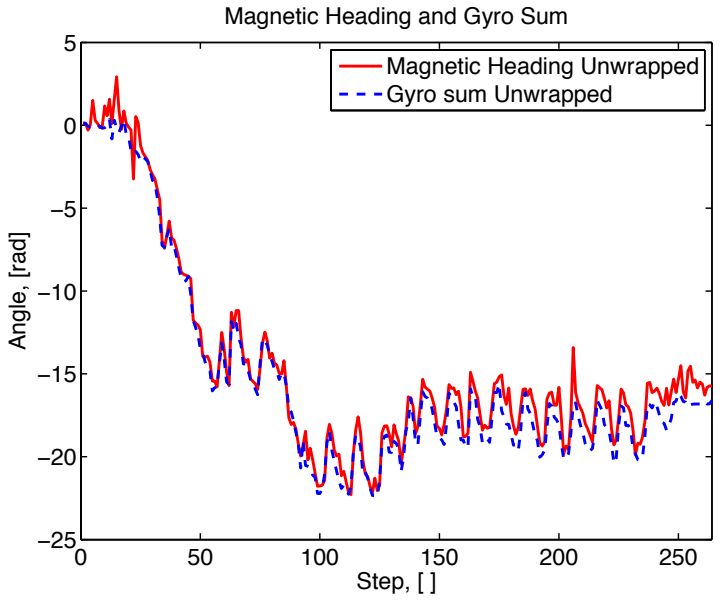


*(b) Experiment 2. The magnetic disturbances are small and primarily present in the beginning and in the end. The gyro drift is quite small considering the number of turns, and is in the end just over a radian.*

***Figure 3:*** *Magnetic heading in solid red and heading gyro summation heading in dashed blue for both experiments. The magnetic heading is quite often disturbed and the gyro based estimation drifts. Still, the shapes of the two vectors are similar.*

In the offices experiment the heading vectors are very similar besides from a few magnetic disturbances in the beginning and the end and a slight drift in the gyro heading.

## 5.1 Principles for Utilizing Magnetic Heading in an INS

As should be clear from Figure 3, the heading can be well estimated by adjusting the gain and bias in the gyro sum, using good segments of data from the magnetometer. This problem is fundamentally a smoothing problem, since we cannot with high confidence determine if the magnetic heading is reliable until we have more data. Introducing smoothing into the INS framework is not a principal problem. However, the application requires a real-time implementation, which is not feasible here.

The heading smoothing problem itself can be solved recursively with little processing requirements, as will be demonstrated below. To utilize this smoothed heading estimate, we propose to utilize the already available information in a simple and standard dead-reckoning system. That is, we propose the following three step algorithm for indoor positioning:

- Estimate $\mathbf{x}_{0:k}$ using the inertial navigation system as described in Sections 2–4. Extract trajectory $(\hat{p}^x_{0:\bar{k}}, \hat{p}^y_{0:\bar{k}})$ and headings $\hat{\psi}_{0:\bar{k}}$.
- Estimate $\hat{\hat{\psi}}_{0:\bar{k}}$ using $\hat{\psi}_{0:\bar{k}}$ and $y^{\psi}_{0:\bar{k}}$.
- Produce a new trajectory $(\hat{\hat{p}}^x_{0:\bar{k}}, \hat{\hat{p}}^y_{0:\bar{k}})$ using $\hat{\hat{\psi}}_{0:\bar{k}}$.

In this way a smoothing estimate of the heading vector is acquired that is used to update the user trajectory. As a result a smooth and accurate user trajectory estimate is produced.

## 5.2 Magnetic Disturbance Detection

The magnetic disturbance detection is based on an HMM filter. For test statistic, it studies the difference between the magnetometer heading measurements and the gyro measurements

$$\lambda^m_{\bar{k}} = y^m_{\bar{k}} - y^m_{\bar{k}-1} - u^{\psi}_{\bar{k}} \tag{14}$$

where $u^{\psi}_{\bar{k}} = \hat{\psi}_{\bar{k}} - \hat{\psi}_{\bar{k}-1}$. If the magnetic heading is undisturbed, $p(\lambda^m_{\bar{k}}|H_0) = \mathcal{N}(0, 2\sigma^2_m + \sigma^2_{\psi,\bar{k}})$. For the disturbed case the distribution has been approximated as $p(\lambda^m_{\bar{k}}|H_1) = \mathcal{N}(0, \sigma^{2,}_{d,\bar{k}})$ where $\sigma^2_{d,\bar{k}} \gg 2\sigma^2_m + \sigma^2_{\psi,\bar{k}}$.

The mode transition probability used was

$$\Pi^m = \begin{pmatrix} 0.95 & 0.05 \\ 0.01 & 0.99 \end{pmatrix} \tag{15}$$

stating that the probability of remaining undisturbed is 95% while the probability of remaining disturbed is 99%. This means that we consider missing a few

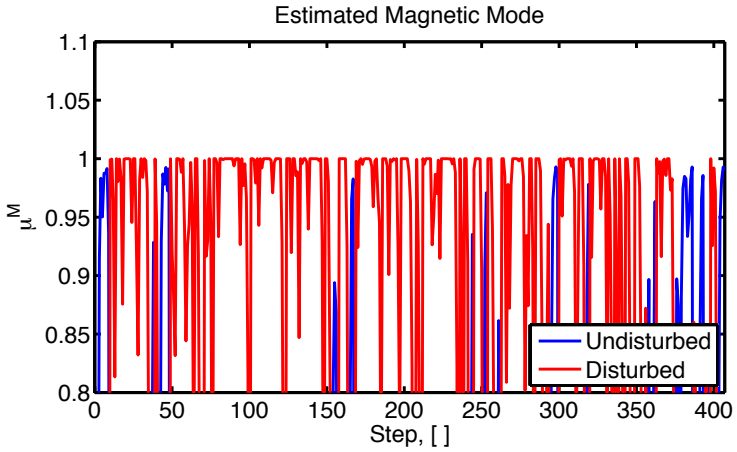


***(a)*** *Experiment 1.*

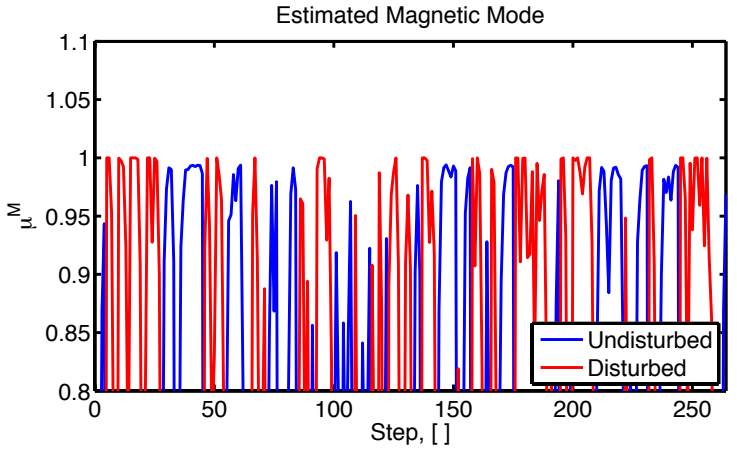


***(b)*** *Experiment 2.*

***Figure 4:*** *Magnetic measurement mode estimation*

good measurements less of a problem than erroneously stating that a disturbed measurement is good.

### Experimental Results

The estimated magnetic mode probabilities for the experiments can be studied in Figure 4. In the first experiment, the magnetic heading is regarded as highly unreliable. Only rarely does the probability of the magnetometer being undisturbed exceed 80%. In the second experiment, roughly half of the magnetic measurements are estimated as undisturbed.

## 5.3 Heading Vector Estimation

The dynamic model of the heading is

$$\psi_{\bar{k}+1} = \psi_{\bar{k}} + u_{\bar{k}}^{\psi} + w_{\bar{k}} \tag{16}$$

where $u_{\bar{k}}^{\psi}$ and $w_{\bar{k}} \sim \mathcal{N}(0, P_{u,\bar{k}})$ are given by the INS.

The measurement equation for the magnetic heading has a possible disturbance $d_{\bar{k}}^{\psi}$. This is modeled as mode dependent measurement noise $e_{\bar{k}}^{\psi} = e(d_{\bar{k}}^{\psi})$.

$$y_{\bar{k}}^{\psi} = \psi_{\bar{k}} + e_{\bar{k}}^{\psi} \tag{17}$$

which can be approximated as $e_{\bar{k}}^{\psi} \sim \mathcal{N}(0, R_m(\mu_{\bar{k}}^m))$. $\mu_{\bar{k}}^m$ has been estimated using an HMM as described above.

If heading was to be estimated recursively using a Kalman filter, the measurement update could be formulated as an unconstrained optimization problem.

$$\begin{aligned} \hat{\psi}_{\bar{k}} = \arg\min_{\psi_{\bar{k}}} \quad & (y_{\bar{k}}^{\psi} - \psi_{\bar{k}})^T R_{m,\bar{k}}^{-1} (y_{\bar{k}}^{\psi} - \psi_{\bar{k}}) + \\ & (\psi_{\bar{k}} - \psi_{\bar{k}-1} - u_{\bar{k}}^{\psi})^T P_{\bar{k}|\bar{k}-1}^{-1} (\psi_{\bar{k}} - \psi_{\bar{k}-1} - u_{\bar{k}}^{\psi}) \end{aligned} \tag{18}$$

where $P_{\bar{k}|\bar{k}-1}$ is the time updated state covariance.

### Estimation on Batch Form

To produce a smoothing estimate $\hat{\psi}_{0:\bar{k}}$, (18) can be written on batch form. Thereby, the entire vector $\hat{\psi}_{0:k}$ is produced in one step using all measurements $y_{0:\bar{k}}^{\psi}$. A prior $\psi^0$ on the initial state $\psi_0$ has been included.

$$\begin{aligned} \hat{\psi}_{0:\bar{k}} = \arg\min_{\psi_{0:\bar{k}}} \quad & (y_{0:\bar{k}}^{\psi} - \psi_{0:\bar{k}})^T R_{m,0:\bar{k}}^{-1} (y_{0:\bar{k}}^{\psi} - \psi_{0:\bar{k}}) + \\ & (\Delta_{\bar{k}} \psi_{0:\bar{k}} - u_{1:\bar{k}}^{\psi})^T Q_{1:\bar{k}}^{-1} (\Delta_{\bar{k}} \psi_{0:\bar{k}} - u_{1:\bar{k}}^{\psi}) + \\ & (\psi_0 - \psi^0)^T P_0^{-1} (\psi_0 - \psi^0) \end{aligned} \tag{19}$$

$P_0$ is the covariance of the prior, $Q_{1:\bar{k}}$ is diagonal and given by the inertial navigation system and

$$\Delta_{\bar{k}} = \begin{pmatrix} 1 & -1 & 0 & 0 \\ 0 & \ddots & \ddots & 0 \\ 0 & 0 & 1 & -1 \end{pmatrix}. \tag{20}$$

When studied, the heading estimates given by the INS is good down to a few error parameters. Figures 3a and 3b indicate that the errors are primarily related to a bias and a slight error in gain.

We therefore assume that the heading can be written as a function of the estimated changes in heading $u_{0:\bar{k}}^{\psi}$, corrected by a bias and a gain term, $b$ and $g$,

respectively, and given a new initial heading.

$$\begin{aligned}\psi_{\bar{k}} &= \psi_0 + \sum_{i=1}^{\bar{k}} (g u_i^{\psi} + b) \\ &= \psi_0 + g\bar{u}_{\bar{k}}^{\psi} + \bar{k}b \end{aligned} \tag{21}$$

where $\bar{u}_{\bar{k}}^{\psi} = \sum_{i=1}^{\bar{k}} u_i^{\psi}$. Putting (21) into (19) gives the heading estimation

$$\begin{aligned}\min_{\psi_0,g,b} \quad & (y_{0:\bar{k}}^{\psi} - \psi_0\mathbf{1} - g\bar{u}_{1:\bar{k}}^{\psi} - b\kappa_{0:\bar{k}})^T R_{m,0:\bar{k}}^{-1}(y_{0:\bar{k}}^{\psi} - \psi_0\mathbf{1} - g\bar{u}_{1:\bar{k}}^{\psi} - b\kappa_{0:\bar{k}}) \\ & \Big((g-1)\bar{u}_{1:\bar{k}}^{\psi} + b\mathbf{1}\Big)^T Q_{0:\bar{k}}^{-1}\Big((g-1)\bar{u}_{1:\bar{k}}^{\psi} + b\mathbf{1}\Big) + \\ & (\psi_0 - \psi^0)^T P_0^{-1}(\psi_0 - \psi^0)\end{aligned} \tag{22}$$

where $\kappa_{0:\bar{k}} = (0\ 1\ 2\ \cdots\ \bar{k})^T$ and $\mathbf{1} = (1\ 1\ \cdots\ 1)$.

The middle term is the cost of deviating from the nominal values $g^0 = 1$ and $b^0 = 0$. It therefore functions as a prior on $b$ and $g$, but one where the cost balance is maintained. Since the dimension of $y_{0:\bar{k}}^{\psi}$ keeps growing, the cost associated to the first term will get bigger after every step. Since the middle term is depending on $\bar{u}_{0:\bar{k}}^{\psi}$ the cost of it also grows with each step, why the balance between the terms are maintained.

The sought heading $\hat{\hat{\psi}}_i$ can now be calculated using (21).

This estimation system is based on a couple of assumptions and simplifications. For example, we assume that the error model (21) is correct. This is not entirely true since the time between two steps is not constant why $b$ is not completely constant. Figure 3 suggests though that the drift in heading seems to be fairly constant after all, making the approximation reasonable.

#### Implementation Issues

We note the change in heading $u_{\bar{k}}^{\psi}$ from the inertial navigation system and sum up all these changes to get an unwrapped vector of gyro based heading estimates, $u_{0:\bar{k}}^{\psi}$. This vector is plotted in red in Figures 3a and 3b.

The vector of magnetic yaw measurements $y_{0:\bar{k}}^{\psi}$ also needs to be unwrapped, making it look like the blue vectors in Figures 3a and 3b. This unwrapping step is difficult but crucial and it has to be done right for the heading estimation to work. This procedure is described in more detail in Callmer et al. (2013b). The heading estimation system described above is very similar to the system suggested in Callmer et al. (2013b).

The measurement covariance matrix $R_{m,0:\bar{k}}$ should not really be diagonal since it depends on the estimated mode probabilities $\mu_{0:\bar{k}}^{m}$ that are dependent over time. In the implementation, $R_{m,0:\bar{k}}$ has been approximated as a diagonal matrix to simplify the computations.

The computational complexity of the problem is linear with time. To solve (22), the cost function is differentiated with respect to $\psi_0$, $b$ and $g$ and is put equal to zero. Solving (22) then boils down to solving

$$\Omega\theta = \gamma \tag{23}$$

where $\theta = \begin{pmatrix} \psi_0 & g & b \end{pmatrix}^T$. It can be solved very cheaply using Cholesky decomposition. To avoid recomputing $\Omega$ and $\gamma$ after each new step, they can be updated using Cholesky rank-one updates. In the end, solving (22) can be done extremely efficiently. More on this in Callmer et al. (2013b).

#### Experimental Results

The heading estimation system has been used to estimate the headings in the experiments described earlier. For Experiment 1, the estimated heading is shown in Figure 5a. It very well reflects the main outline of the magnetic heading vector while maintaining the basic shape of the gyro sum vector. In Experiment 2, the corridor experiment, the bias in the heading estimate has now been removed, Figure 5b.

## 5.4 Trajectory Postprocessing using Heading

The user positions $\hat{p}^x_{0:\bar{k}}$ and $\hat{p}^y_{0:\bar{k}}$ can be written as a series of steps from the origin. A position $\hat{p}_{i+1} = (\hat{p}^x_{i+1}, \hat{p}^y_{i+1})$ is based on the previous position $\hat{p}_i = (\hat{p}^x_i, \hat{p}^y_i)$, the step taken $(\Delta x_{i+1}, \Delta y_{i+1})$ and the previous direction $\hat{\psi}_i$. Each step is taken in the local coordinate system of the previous step, see Figure 6. $(\hat{p}^x_{i+1}, \hat{p}^y_{i+1})$ subsequently becomes

$$\begin{pmatrix} \hat{p}^x_{i+1} \\ \hat{p}^y_{i+1} \end{pmatrix} = \begin{pmatrix} \cos(\hat{\psi}_i) & -\sin(\hat{\psi}_i) \\ \sin(\hat{\psi}_i) & \cos(\hat{\psi}_i) \end{pmatrix} \begin{pmatrix} \Delta x_{i+1} \\ \Delta y_{i+1} \end{pmatrix} + \begin{pmatrix} \hat{p}^x_i \\ \hat{p}^y_i \end{pmatrix} \tag{24}$$

Replacing $\hat{\psi}_{0:\bar{k}}$ with the newly estimated $\hat{\hat{\psi}}_{0:\bar{k}}$, a new user trajectory $(\hat{\hat{p}}^x_{0:\bar{k}}, \hat{\hat{p}}^y_{0:\bar{k}})$ can be calculated using (24).

## 5.5 Magnetic Heading Utilization Discussion

Another option of how to utilize the magnetic heading measurements is to include them in the inertial navigation framework. The measurement update for a stand still can be extended to also include the magnetic measurements. Magnetic disturbances can for example be discarded by studying the filter innovation, or a more elaborate scheme such as a Kalman filter bank as in Callmer et al. (2013c), can be used. Extensive experimenting indicates though that such an approach does not work well in practice. The filtering estimates of position are very sensitive to minor changes in heading making the trajectory appear extremely 'nervous', at best. To get a smooth trajectory estimate, one would have to reestimate the entire trajectory after each magnetic heading update using smoothing. In the end, such an approach will produce results similar to the ones presented here, only immensely more computationally complex.

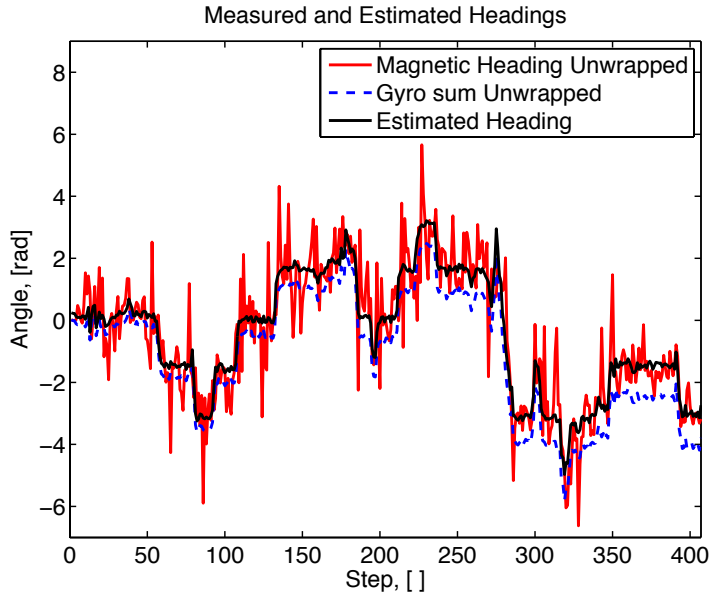


***(a)*** *Experiment 1. The estimated heading is basically centered around the magnetic heading since the magnetic disturbances seem to be distributed quite evenly throughout the experiment.*

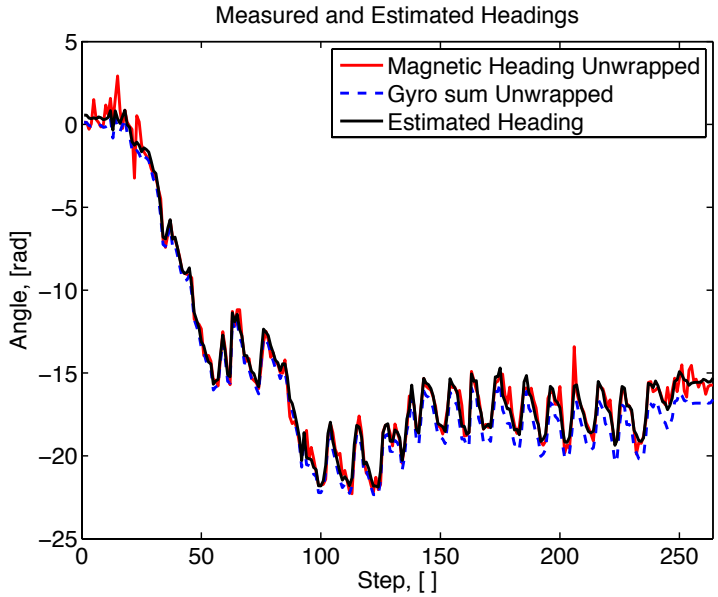


***(b)*** *Experiment 2. The heading estimate has been adjusted to the magnetic heading vector and is now aligned well also in the early sections of the dataset.*

***Figure 5:*** *Heading estimation. Magnetic heading in red, inertial navigation system based heading in dashed blue and estimated heading in black.*

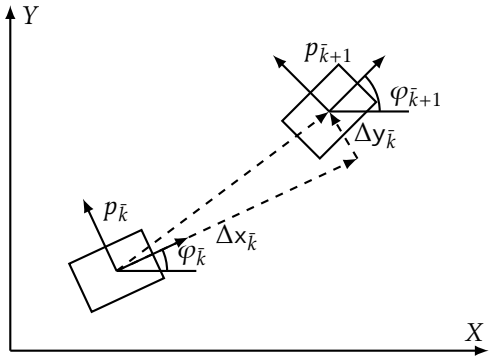


***Figure 6:*** *Coordinate systems used to describe position and movement. User position at time $\bar{k}$ is $p_{\bar{k}} = (p^x_{\bar{k}}, p^y_{\bar{k}})$ and orientation is $\psi_{\bar{k}}$. A step $(x_{\bar{k}+1}, y_{\bar{k}+1})$ is taken in the local coordinate system from the previous time.*

# 6 Experimental Results

We will now show the experimental positioning results using the heading estimation and trajectory adjustment system. For performance evaluation, the final results have been overlaid on floor plans of the building.

## 6.1 Experiment 1

The estimated trajectory $(\hat{p}^x_{0:\bar{k}}, \hat{p}^y_{0:\bar{k}})$ for the multiple floors experiment is shown in Figure 7. Since the user has walked on multiple floors, both floors are plotted. The trajectory has been color coded based on the estimated elevation. If the estimated elevation was below 2.5 meters, the track was blue, while all paths above 2.5 meters were made red. For this experiment this means that the ground floor results are blue and the second floor paths are red. Note how the color switches coincide with the stairs in the floor plan.

The start and end position was the same making the final position error about 5 meters, barely 1% of distance travelled. For the left half of the trajectory, the estimate is a few meters off, pulled a bit down. A significant part of this error was created while walking the stairs indicated by the arrows in Figure 7. An error of a few meters was accumulated that could not be corrected. Studied more closely, almost all floor switches seem to induce a bit of position drift. Apparently, the stairs are causing problems for the inertial navigation system.

## 6.2 Experiment 2, Corridor

The second experiment is the corridor dataset. The estimated trajectory has been overlaid on a floor plan in Figure 8. The two paths along the corridors are well aligned and the offices are perpendicular to the corridor. The final error in position is about 2 meters or less than 1%. Slight errors in position along the

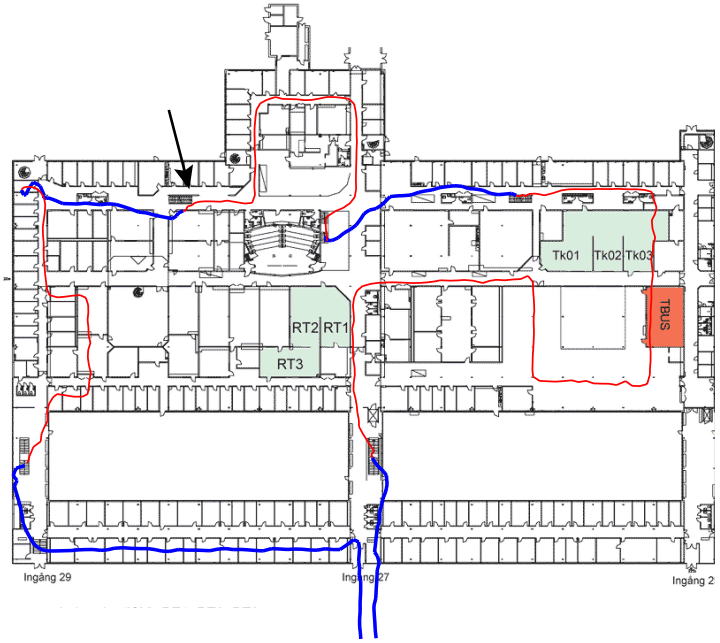


***(a)*** *Ground floor. Blue tracks are paths on ground floor.*

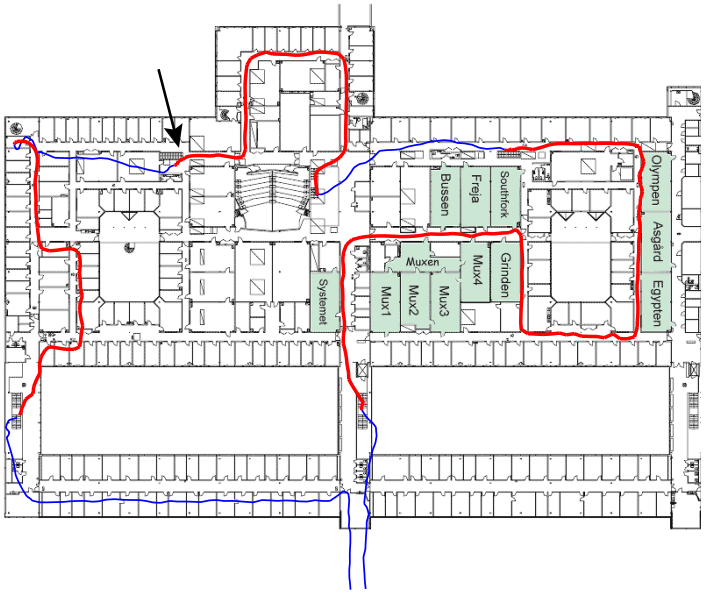


***(b)*** *Second floor. Red tracks are paths on second floor.*

***Figure 7:*** *Experimental positioning results overlaid on a floor plan of the building. The estimate is at the most about 5 meters off. A majority of the errors where accumulated while walking up or down the stairs. The stair indicated by the arrow caused a large part of the total positioning error.*

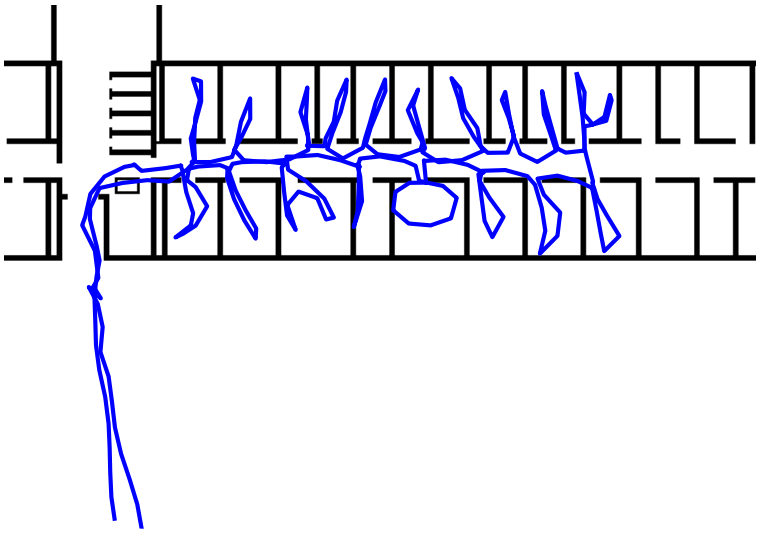

***Figure 8:*** *Estimated trajectory of the corridor experiment overlaid on a layout of the corridor.*

trajectory can be observed when the results are studied closely, but no errors are larger than 1-2 meters. The door of each office is positioned accurately. The tilt and shape of the individual path walked in each office is actually depending on the furnishing in each respective office.

## 7 Conclusions

There are many successful results reported in literature on positioning based on foot-mounted IMMU, with relative errors as small as 1%. The zero-velocity update is the key to get a small error in walked distance. However, the weak link seems to be the heading estimate. Either the authors have evaluated their system outdoors to avoid magnetic disturbances, have used building information explicitly (maps) or implicitly (cornering is restricted to ±90 degrees) or required that previously visited areas are revisited. We propose a novel three step algorithm for robust positioning in real-time:

1) A Kalman filter framework based on a 3D navigation model, where the discrete modes of stand still are estimated with an HMM filter, triggering a soft measurement update in the KF. 2) A separate smoothing filter for heading based on the KF output and magnetometers. 3) A simple dead-reckoning filter to correct the position from the KF with the robust heading estimate.

Our experiments demonstrate an improvement from almost 10% relative error

and roughly $\pi/4$ radian heading error in the final position from the KF, to less than 1% position error with the three-step algorithm with virtually no heading error.

## Acknowledgments

The authors would like to thank Jouni Rantakokko and FOI for lending us the IMMUused in the experiments.

This work was supported by CADICS, a Linnaeus center funded by the Swedish Research Council, and by the Excellence Center at Linköping-Lund in Information Technology (ELLIIT).

## Bibliography

K. Abdulrahim, C. Hide, T. Moore, and C. Hill. Aiding low cost inertial navigation with building heading for pedestrian navigation. *Journal of Navigation*, 64: 219–4233, 2011.

M. Angermann and P. Robertson. FootSLAM: Pedestrian simultaneous localization and mapping without extroceptive sensors – hitchiking on human perception and cognition. *Proceedings of the IEEE*, 100(13):1840–1848, 2012.

E. Bachmann, X. Yun, and A. Brumfield. Investigation of the effects of magnetic variations on inertial/magnetic orientation sensors. *IEEE Robotics and Automation Magazine*, 14(3):76 – 87, 2007.

E. Bachmann, J. Calusdian, E. Hodgson, and X. Yun. In situ heading drift correction for human position tracking using foot-mounted inertial/magnetic sensors. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012.

P. Bahl and V.N Padmanabhan. RADAR: an in-building rf-based user location and tracking system. In *Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies INFOCOM*, 2000.

S. Beauregard. Omnidirectional pedestrian navigation for first responders. In *Proc. of the 4th Workshop on Positioning, Navigation and Communication, WPNC07*, Hannover, Germany, 2007.

Ö. Bebek, M.A. Suster, S. Rajgopal, M.J. Fu, X. Huang, M.C. Çavuşoğlu, D.J. Young, M. Mehregany, A.J. van den Bogert, and C.H. Mastrangelo. Personal navigation via high-resulution gait-corrected inertial measurement units. *IEEE Transactions on Instrumentation and Measurement*, 59(11):3018–3027, 2010.

J. Borenstein and L. Ojeda. Heuristic drift elimination for personnel tracking systems. *Journal of Navigation*, 63(3):591 – 606, 2010.

J. Borenstein, L. Ojeda, and S. Kwanmuang. Heuristic reduction of gyro drift in imu-based personnel tracking systems. In *Proceedings of SPIE, the International Society for Optical Engineering*, 2009.

J. Callmer. Topics in localization and mapping. Linköping University, Sweden, 2011. Licentiate Thesis No 1489.

J. Callmer, D. Törnqvist, and F. Gustafsson. Probabilistic stand still detection using foot mounted IMU. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

J. Callmer, D. Törnqvist, and F. Gustafsson. An inertial navigation framework for indoor positioning with robust heading. *IEEE Transactions on Instrumentation and Measurement*, 2013a. Submitted.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors

using convex optimization. In *International Conference on Information Fusion*, 2013b. Submitted.

J. Callmer, D. Törnqvist, and F. Gustafsson. Robust heading estimation indoors. *IEEE Transactions on Signal Processing*, 2013c. Submitted.

W. H. K. de Vries, H. E. J. Veeger, C. T. M. Baten, and F. C. T. van der Helm. Magnetic distortion in motion labs, implications for validating inertial magnetic sensors. *Gait Posture*, 29(4):1786 – 1792, 2009.

E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

S. Godha, G. Lachapelle, and M. E. Cannon. Integrated GPS/INS system for pedestrian navigation in signal degraded environment. In *Proc. of ION GNSS*, 2006.

T. Harada, H. Uchino, T. Mori, and T. Sato. Portable absolute orientation estimation device with wireless network under acceleration situation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

V. Honkavirta, T. Perala, S. Ali-Loytty, and R. Piché. A comparative survey of WLAN location fingerprinting methods. In *Workshop on Positioning, Navigation and Communication (WPNC 2009)*, 2009.

A.R. Jiménez, F. Seco, J.C. Prieto, and J. Guevara. Indoor pedestrian navigation using an ins/ekf framework for yaw drift reduction and a foot-mounted imu. In *7th Workshop on Positioning, Navigation and Communication, WPNC*, 2010.

C. W. Kang and C. G. Park. Improvement of attitude estimation using hidden markov model classification. In *AIAA Infotech@Aerospace 2010 Conference and Exhibit*, 2010.

B. Krach and P. Robertson. Cascaded estimation architecture for integration of foot-mounted inertial sensors. In *Proc of the IEEE/ION Position, Location and Navigation Symposium*, 2008.

J.B. Kuipers. *Quaternions and Rotation Sequences*. Princeton University Press, 1999.

J. K. Lee and E. J. Park. Minimum-order kalman filter with vector selector for accurate estimation of human body orientation. *IEEE Transactions on Robotics*, 25(5):1196 – 1201, 2009.

B. Li, J. Salter, A.G. Dempster, and C. Rizos. Indoor positioning techniques based on wireless LAN. Technical report, School of Surveying and Spatial Information Systems, UNSW, Sydney, Australia, 2006.

J.-O. Nilsson, I. Skog, and P. Händel. A note on the limitations of ZUPTs and the implications on sensor error modeling. In *International Conference in Indoor Positioning and Indoor Navigation*, 2012.

L. Ojeda and J. Borenstein. Non-GPS navigation for security personnel and first responders. *Journal of Navigation*, 60(3):391–407, 2007.

J. Rantakokko, J. Rydell, P. Strömbäck, P. Händel, J. Callmer, D. Törnqvist, F. Gustafsson, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: multisensor systems and cooperative localization. *Wireless Communications, IEEE*, 18(2):10–18, 2011.

V. Renaudin, O. Yalak, P. Tomé, and B. Merminod. Indoor navigation of emergency agents. *European Journal of Navigation*, 5(3), 2007.

D. Roetenberg, H. J Luinge, C. T. M. Baten, and P. H. Veltink. Compensation of magnetic disturbances improves inertial and magnetic sensing of human body segment orientation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3):395 – 405, 2005.

A. M. Sabatini. Estimating three-dimensional orientation of human body parts by inertial/magnetic sensing. *Sensors*, 11(2):1489 – 1525, 2011.

J. Seitz, T. Vaupel, J. Jahn, S. Meyer, J.G. Boronat, and J. Thielecke. A hidden markov model for urban navigation based on fingerprinting and pedestrian dead reckoning. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

I. Skog, P. Händel, J.-O. Nilsson, and J. Rantakokko. Zero-velocity detection – an algorithm evaluation. *IEEE Transactions on Biomedical Engineering*, 57, 2010.

R. Stirling, J. Collin, K. Fyfe, and G. Lachapelle. An innovative shoe-mounted pedestrian navigation system. In *Proceedings of European Navigation Conference GNSS*, 2003.

U. Walder, T. Bernoulli, and T. Wießflecker. An indoor positioning system for improved action force command and disaster management. In *Proc. of the 6th Int iscram Cont*, 2009.

Widyawan, M. Klepal, and S. Beauregard. A backtracking particle filter for fusing building plans with PDR displacement estimates. In *Positioning, Navigation and Communication, 2008. WPNC 2008. 5th Workshop on*, pages 207 –212, march 2008.

O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing (UbiComp)*, 2008.

O. Woodman and R. Harle. Rf-based initialisation for inertial pedestrian tracking. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

X. Yun, J. Calusdian, E. Bachmann, and R.B. McGhee. Estimation of human foot motion during normal walking using inertial and magnetic sensor measurements. *IEEE Transactions on Instrumentation and Measurement*, 61(7):2059–2072, 2012.

# Paper D

# RADAR SLAM using Visual Features

*Authors:* Jonas Callmer, David Törnqvist, Henrik Svensson, Pelle Carlbom and Fredrik Gustafsson

*Edited version of the paper:*

J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011.

# RADAR SLAM using Visual Features

Jonas Callmer*, David Törnqvist*, Henrik Svensson**,
Pelle Carlbom† and Fredrik Gustafsson*

*Dept. of Electrical Engineering,
Linköping University,
SE–581 83 Linköping, Sweden
callmer@isy.liu.se

**Nira Dynamics,
Linköping, Sweden

†Saab Dynamics
Linköping, Sweden

**Abstract**

A vessel navigating in a critical environment such as an archipelago, requires very accurate movement estimates. Intentional or unintentional jamming makes GPS unreliable as the only source of information and an additional independent supporting navigation system should be used. In this paper we suggest estimating the vessel movements using a sequence of RADAR images from the preexisting body-fixed radar. Island landmarks in the RADAR scans are tracked between multiple scans using visual features. This provides information not only about the position of the vessel but also of its course and velocity. We present here a navigation framework that requires no additional hardware than the already existing naval RADAR sensor. Experiments show that visual RADAR features can be used to accurately estimate the vessel trajectory over an extensive data set.

## 1 Introduction

In autonomous robotics there is a need to accurately estimate the movements of a vehicle. A simple movement sensor like a wheel encoder on a ground robot or a pit log on a vessel will under ideal circumstances provide quite accurate movement measurements. Unfortunately they are sensitive to disturbances. For example, wheel slip due to a wet surface will be interpreted incorrectly by a wheel encoder and strong currents will not be correctly registered by the pit log why a position estimate based solely on these sensors will drift off. In applications like autonomous robotics the movement accuracy needs to be high why other redundant movement measurement methods are required.

A common approach is to study the surroundings and see how they change over time. By relating the measurements of the environment $k$ seconds ago to the present ones, a measurement of the vehicle translation and rotation during this time interval can be obtained. A system like this complements the movement

***Figure 1:*** *The high speed patrol boat type used for the data acquisition. Note the backwash created by the jet propulsion system. Courtesy of Dockstavarvet AB.*

sensor and enhances the positioning accuracy.

Most outdoor navigation systems such as surface vessels use global navigation satellite systems (GNSS) such as the Global Positioning System (GPS) to measure their position. These signals are weak making them very vulnerable to intentional or unintentional jamming Volpe (2001); GNS (2001); Grant et al. (2009). A supporting positioning system that is redundant of the satellite signals is therefore necessary. By estimating the vessel movements using the surroundings, a mean of measuring the reliability of the GPS system is provided. The movement estimates can also be used during a GPS outage providing accurate position and movement estimates over a limited period of time. This support system could aid the crew in critical situations during a GPS outage, avoiding costly PR disasters such as running aground.

For land based vehicles or surface vessels, three main sensor types exist that can measure the environment: cameras, laser range sensors and radar sensors. Cameras are very rich in information and has a long reach but are sensitive to light and weather conditions. Laser range sensors provide robust and accurate range measurements but also they are very sensitive to weather conditions. The radar signal is usually the least informative signal of the three and is also quite sensitive to what the signals reflect against. On the other hand, the radar sensor works pretty much equally well in all weather conditions.

In this paper, radar scan matching to estimate relative movements is studied. The idea is to use the RADAR as an imagery sensor, and apply computer vision algorithms to detect landmarks of opportunity. Landmarks that occur during consecutive RADAR scans are then used for visual odometry, that gives speed, relative position and relative course. The main motivation for using visual features to match RADAR scans instead of trying to align the RADAR scans is that visual

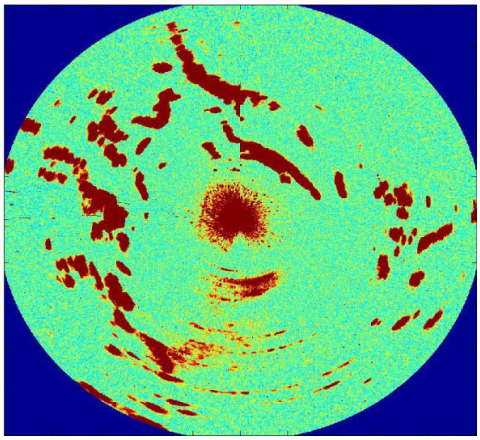

***Figure 2:*** *Typical RADAR image showing the islands surrounding the vessel. The RADAR disturbances close to the vessel are caused by the vessel and the waves. Behind the vessel (lower part of the image) the striped shaped disturbances are the result of backwashes reflecting the RADAR pulses.*

features are easily matched despite large translational and rotational differences, which is more difficult using other scan matching techniques. The landmarks can optionally be saved in a map format which can be used to recognize areas that have been visited before. That is, a by-product of the robust navigation solution is a mapping and exploration system.

Our application example is based on a military patrol boat, Figure 1, that often maneuvers close to the shore in high speeds, at night, without visual aid in situations where GPS jamming or spoofing cannot be excluded. As the results will show, we are able to navigate in a complex archipelago using only the radar, and get a map that is very close to ground truth.

To provide a complete backup system for GPS, global reference measurements are necessary to eliminate the long term drift. The surface navigation system in Dalin and Mahl (2007); Karlsson and Gustafsson (2006), assumed that an accurate sea chart is available. The idea was to apply map matching between the RADAR image and the sea chart, and the particle filter was used for this mapping. Unfortunately, commercial sea charts still contain rather large absolute errors of the shore, see Volpe (2001); GNS (2001), which makes them less useful in blind navigation with critical maneuvers without visual feedback.

The RADAR used in these experiments measures the distances to land areas using 1024 samples in each direction and a full revolution is comprised of roughly 2000 directions. Each scan has a radius of about 5 km giving a range resolution of roughly 5 meters. These measurements are used to create a RADAR image by translating the range and bearing measurements into Cartesian coordinates. An example of the resulting image is shown in Figure 2.

The RADAR image gives a birds eye view of the surrounding islands and by tracking these islands, information about how the vessel is moving is obtained. We use the Scale-Invariant Feature Transform (SIFT) Lowe (1999) to extract trackable features from the RADAR image which are subsequently matched with features from later scans. These features are shown to be distinct and stable enough to be used for island tracking. Other feature detectors like Speeded Up Robust Features (SURF) Bay et al. (2006) could equally well have been used. When these features are tracked using a filter, estimates of the vessel movements are obtained that over time give an accurate trajectory estimate.

The outline is as follows; Section 2 gives a overview of the related work followed by a theoretical filtering framework in Section 3. In Section 4, the performance of SIFT is evaluated on RADAR images and the trajectory estimation performance on experimental data is given in Section 5. The paper then ends in Section 6 with conclusions and suggested future work.

## 2 Background and Relation to SLAM

The approach in this contribution is known as the Simultaneous Localization And Mapping (SLAM) problem. Today, SLAM is a fairly well studied problem with solutions that are reaching some level of maturity Durrant-Whyte and Bailey (2006); Bailey and Durrant-Whyte (2006). SLAM has been performed in a wide variety of environments such as indoors Newman and Ho (2005), in urban Bosse and Zlot (2008); Granström et al. (2009); Rouveure et al. (2009); Checchin et al. (2009) and rural areas Ramos et al. (2007b); Rouveure et al. (2009), underwater Eustice et al. (2005); Mahon et al. (2008) and in the air Bryson and Sukkarieh (2005) and the platform is usually equipped with a multitude of sensors such as lasers, cameras, inertial measurement units, wheel encoders, etc. In this work we will use only the RADAR sensor of a naval vessel to perform SLAM in a maritime environment. The data used was recorded in the Stockholm archipelago by Saab Bofors Dynamics Carlbom (2005).

Radars have been used for a long time to estimate movements, for example in the early experiments by Clark and Durrant-Whyte (1998). Radar reflecting beacons in known positions were tracked using a millimeter RADAR and this was shown to improve the movement estimates. Shortly after, Clark and Dissanayake (1999) extended the work by tracking natural features instead of beacons.

Thereafter, laser range sensors became more popular since they are more reliable, giving a range measurement in all directions. The problem of estimating the vehicle movements became a problem of range scan alignment. This was studied among others in Feng and Milios (1994); Lu and Milios (1997); Chen and Medioni (1992); Ramos et al. (2007a).

The advantages of the radar, such as its ability to function in all weather conditions, have though resulted in it making a comeback. Lately, microwave radars have been used in SLAM experiments but now using a landmark free approach.

In Rouveure et al. (2009), SLAM was performed in both urban and rural areas by aligning the latest RADAR scan with the RADAR map using 3D correlations to estimate the relative movements of the vehicle. The RADAR map was constructed by consecutively adding the latest aligned RADAR scan to the previous scans. Checchin et al. (2009) performed SLAM in an urban scenario by estimating the rotation and translation of the robot over a sequence of scans using the Fourier-Mellin Transform. It can match images that are translated, rotated and scaled and can therefore be used to align RADAR scans Chen et al. (1994). Chandran and Newman (2006) jointly estimated the RADAR map and the vehicle trajectory by maximizing the quality of the map as a function of a motion parametrization.

Millimeter wave radars have also become more commonplace in some segments of the automotive industry and the number of applications for them are growing. For example, the road curvature has been estimated using the radar reflections which will be used in future systems in collision warning and collision avoidance Tsang et al. (2006); Lundquist (2009).

The problem of RADAR alignment is also present in meteorology where space RADAR and ground RADAR observations are aligned to get a more complete picture of the weather in Bolen and Chandrasekar (2003). The scans are aligned by dividing them into smaller volumes that are matched by their respective precipitation intensities.

Visual features like SIFT or SURF have been used in camera based SLAM many times before. Sometimes the features were used to estimate relative movements Se et al. (2002); Eustice et al. (2005); Jensfelt et al. (2006), other times they were used to detect loop closures Cummins and Newman (2007); Newman and Ho (2005); Callmer et al. (2008); Mahon et al. (2008).

The combination of RADAR and SIFT has previously been explored by in Li et al. (2008), where Synthetic Aperture Radar measurements were co-registered using matched SIFT features. Radar scan matching using SIFT was also suggested in the short papers Schikora and Romba (2009); Essen et al. (2010). A system with parallel stationary ground radars are discussed and SIFT feature matching is suggested as a way to estimate the constant overlaps between the scans. No RADAR scans ever seem to be matched in those papers though. To the best of the authors knowledge, this is the first time visual features have been used to estimate the rotational and translational differences between RADAR images

# 3 Theoretical Framework

All vessel movements are estimated relative a global position. The positions of the tracked landmarks are not measured globally but relative to the vessel. Therefore two coordinate systems are used, one global for positioning the vessel and all the landmarks and one local relating the measured feature positions to the vessel. Figure 3 shows the local and global coordinate systems, the vessel and a landmark $m$.

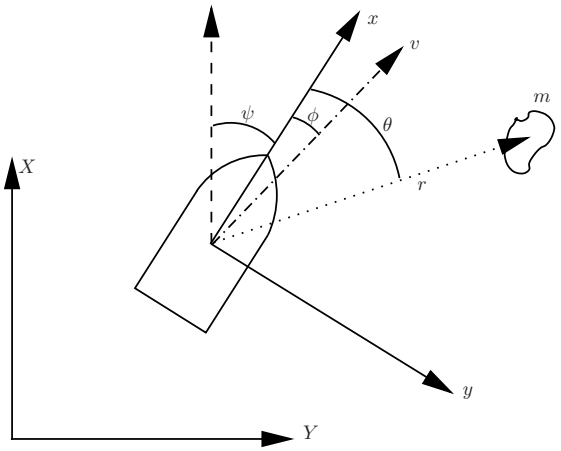


***Figure 3:*** *The global $(X, Y)$ and the local boat fixed $(x, y)$ coordinate systems, the course $\psi$ and the crab angle $\phi$ giving difference between the course and velocity vector. The vessel and an island landmark $m$ are also depicted as is the measured bearing $\theta$ and range $r$ of the landmark relative to the vessel.*

The variables needed for visual odometry are summarized in Table 1.

## 3.1 Detection Model

Each RADAR scan has a radius of about 5 km with a range resolution of 5 meters and the antenna revolution takes about 1.5 s.

If a landmark is detected at time $t$, the RADAR provides a range, $r_t$, and bearing, $\theta_t$, measurement to the island landmark $i$ as

$$\mathrm{y}_t^i = \begin{pmatrix} r_t^i \\ \theta_t^i \end{pmatrix} + e_t^i \tag{1}$$

where $e_t^i$ is independent Gaussian noise. These echos are transformed into a RADAR image using polar to rectangular coordinates conversion and the result is shown in Figure 2. Figures 1 and 2 also show that the forward and sideways facing parts of the scans are the most useful ones for feature tracking. This is due to the significant backwash created by the jet propulsion system of the vessel, which is observed along the vessel trajectory in Figure 1. This backwash disturbs the RADAR measurements by reflecting the RADAR pulse, resulting in the stripe shaped disturbances behind the vessel in Figure 2.

SIFT is today a well established standard method to extract and match features from one image to features extracted from a different image covering the same scene. It is a rotation and affine invariant Harris point extractor that uses a difference-of-Gaussian function to determine scale. Harris points are in turn regions in the image where the gradients of the image are large, making them prone

***Table 1:*** *Summary of notation*

| Parameter | Description |
|---|---|
| $(X, Y)$ | Global position |
| $(x, y)$ | Local coordinate system with $x$ aligned with the stem |
| $v$ | Velocity |
| $\psi$ | Course, defined as the angle between the global $X$ axis and the local $x$ axis, as would be shown by a compass |
| $\omega$ | Turn rate, defined as $\dot{\psi}$ |
| $\phi$ | Difference between course and velocity vector ("crab" angle), which is mainly due to streams and wind |
| $(m_X^i, m_Y^i)$ | Global position of landmark $i$ |
| $r^i$ | Range from the strap-down body-fixed RADAR to landmark $i$ |
| $\theta^i$ | Bearing from the RADAR to landmark $i$ |

to stand out also in other images of the same area. For region description, SIFT uses gradient histograms in 16 subspaces around the point of interest.

In this work, SIFT is used to extract and match features from RADAR images. By tracking the SIFT features over a sequence of RADAR images, information about how the vessel is moving is obtained.

## 3.2 Measurement Model

Once a feature has been matched between two scans, the position of the feature is used as a measurement to update the filter. Since the features are matched in Cartesian image coordinates, the straightforward way would be to use the pixel coordinates themselves as a measurement. After having first converted the pixel coordinates of the landmark to coordinates in the local coordinate system, the Cartesian feature coordinates are now related to the vessel states as

$$\bar{y}_t^i = \begin{pmatrix} y_{x,t}^i \\ y_{y,t}^i \end{pmatrix} + \bar{e}_t^i = R(\psi_t) \begin{pmatrix} m_{X,t}^i - X_t \\ m_{Y,t}^i - Y_t \end{pmatrix} + \begin{pmatrix} \bar{e}_{X,t}^i \\ \bar{e}_{Y,t}^i \end{pmatrix} \tag{2}$$

where $y_{x,t}^i$ is the measured $x$-coordinate of feature $i$ in the local coordinate frame at time $t$ and $R(\psi_t)$ is the rotation matrix between the ship orientation and the global coordinate system. $(X, Y)$ and $(m_X^i, m_Y^i)$ are global vessel position and global position of landmark $i$, respectively.

The problem with this approach is that $\bar{e}_{X,t}$ and $\bar{e}_{Y,t}$ in (2) are dependent since they are both mixtures of the range and bearing uncertainties of the RADAR sensor. These dependencies are also time dependent since the mixtures depend on the bearing of the RADAR sensor. Simply assuming them to be independent will introduce estimation errors.

A better approach is to convert the Cartesian landmark coordinates back to polar coordinates and use these as a measurement

$$y_t^i = \begin{pmatrix} r_t^i \\ \theta_t^i \end{pmatrix} = \begin{pmatrix} \sqrt{(m_{X,t}^i - X_t)^2 + (m_{Y,t}^i - Y_t)^2} \\ \arctan\left(\frac{m_{Y,t}^i - Y_t}{m_{X,t}^i - X_t}\right) - \psi_t \end{pmatrix} + \begin{pmatrix} e_{r,t}^i \\ e_{\theta,t}^i \end{pmatrix}. \tag{3}$$

This approach results in independent noise parameters $e_r \sim \mathcal{N}(0, \sigma_r^2)$ and $e_\theta \sim \mathcal{N}(0, \sigma_\theta^2)$, which better reflect the true range and bearing uncertainties of the range sensor.

## 3.3 Motion Model

The system states describing the vessel movements at time instant $t$ are

$$z_t = \begin{pmatrix} X_t & Y_t & v_t & \psi_t & \omega_t & \phi_t \end{pmatrix}^T \tag{4}$$

where $v$ is the velocity, $\psi$ is the course, $\omega$ is the angular velocity and $\phi_t$ is the crab angle, i.e. the wind and stream induced difference between course and velocity vector (normally small). Due to the size and the speed of the vessel, Figure 1, the crab angle is assumed to be very small throughout the experiments. The system states are more extensively described in Table 1 and are also shown in Figure 3. We will be using a coordinated turn model, though there are many possible motion models available.

When landmarks at unknown positions are tracked to estimate the movements of the vessel, these should be kept in the state vector. If the same landmarks are tracked over a sequence of RADAR scans, a better estimate of the vessel movement is acquired than if they are tracked between just two.

The system states are therefore expanded to also include all landmarks within the field of view to create a visual odometry framework. The new state vector becomes

$$z_t = \begin{pmatrix} X_t & Y_t & v_t & \psi_t & \omega_t & \phi_t & m_{X,t}^k & m_{Y,t}^k & \dots & m_{Y,t}^l \end{pmatrix}^T \tag{5}$$

Only the $l - k + 1$ latest landmarks are within the field of view why only these are kept in the state vector. As the vessel travels on, the landmarks will one by one leave the field of view why they will be removed from the state vector and subsequently replaced by new ones.

When all old landmarks are kept in the state vector even after they have left the field of view, it is a SLAM framework. If an old landmark that left the field of view long ago was rediscovered, this would allow for the whole vessel trajectory to be updated. This is called a loop closure and is one of the key features in SLAM. The SLAM state vector is therefore

$$z_t = \begin{pmatrix} X_t & Y_t & v_t & \psi_t & \omega_t & \phi_t & m_{X,t}^1 & m_{Y,t}^1 \dots \end{pmatrix}^T. \tag{6}$$

A discretized linearization of the coordinated turn model using the SLAM land-

mark augmentation gives

$$\begin{pmatrix} X_{t+\Delta t} \\ Y_{t+\Delta t} \\ v_{t+\Delta t} \\ \psi_{t+\Delta t} \\ \omega_{t+\Delta t} \\ \phi_{t+\Delta t} \\ m^1_{X,t+\Delta t} \\ m^1_{Y,t+\Delta t} \\ \vdots \end{pmatrix} = \begin{pmatrix} X_t + \frac{2v_t}{\omega_t} \sin(\frac{\omega_t \Delta t}{2}) \cos(\psi_t + \phi_t + \frac{\omega_t \Delta t}{2}) \\ Y_t + \frac{2v_t}{\omega_t} \sin(\frac{\omega_t \Delta t}{2}) \sin(\psi_t + \phi_t + \frac{\omega_t \Delta t}{2}) \\ v_t + \nu_{v,t} \\ \psi_t + \omega_t \Delta t \\ \omega_t + \nu_{\omega,t} \\ \phi_t + \nu_{\phi,t} \\ m^1_{X,t} \\ m^1_{Y,t} \\ \vdots \end{pmatrix} \tag{7}$$

where $\Delta t$ is the difference in acquisition time between the two latest matched features. $\nu_v$, $\nu_\omega$ and $\nu_\phi$ are independent Gaussian process noises reflecting the movement uncertainties of the vessel.

## 3.4 Multi-Rate Issues

Having defined a motion model and a measurement model, state estimation is usually straightforwardly implemented using standard algorithms such as the extended Kalman filter (EKF), the unscented Kalman filter (UKF) or the particle filter (PF), see Gustafsson (2010). These filters iterate between a prediction step based on the motion model and a correction step based on the measurement model and the current measurement. The most natural approach is to stack all landmarks from one RADAR revolution into a large measurement vector and then run the filter with $\Delta t = T = 1.5$, where $T$ is the RADAR revolution time (1.5 in our application). There are, however, two non-standard problems here.

The first problem is that all matched features should not be used to update the filter at the same time. Since the vessel is moving while the RADAR is revolving, a shift in the RADAR image is introduced. A vessel speed of 10 m/s will result in a difference in landmark position of about 15 meters from the beginning to the end of the scan. If the vessel is travelling with a constant velocity and course, then all relative changes in feature positions between two scans would be equal, but if the vessel is turning or accelerating, a shift in the relative position change is introduced. This results in different features having different relative movements which will introduce estimation errors if all measurements are used at the same time. We have found the error caused by this batch approach to be quite large. Therefore, the filter should be updated by each matched feature independently.

If independent course and velocity measurements were available they could be used to correct the skewness in the RADAR image. The filter estimates of velocity and course should though not be used for scan correction, since this would create a feedback loop from the estimates to the measurements that can cause filter instability.

Second, if there was one landmark detected in each scan direction, the filter could be updated with the rate $\Delta t = T/N = \frac{1.5}{2000}$ where $N$ is the number of measure-

ments per rotation (2000 in our application). This is not the case though and we are facing a multi-rate problem with irregularly sampled measurements. The measurement model can now conveniently be written as

$$\mathsf{y}_t = \begin{cases} \mathsf{y}_t^i & \text{if landmark } i \text{ is detected at time } t \\ \text{NaN} & \text{otherwise.} \end{cases} \tag{8}$$

Now, any of the filters (EKF, UKF, PF) can be applied at rate $\Delta t = T/N$ using (8), with the understanding that a measurement being NaN simply means that the measurement update is skipped.

## 3.5 Alternative Landmark Free Odometric Framework

The landmark based framework derived above suffers from one apparent shortcoming: the number of features will grow very fast. After only a short time period, thousands of potential landmarks will have been found, causing large overhead computations in the implementation. Either a lot of restrictions must be made on which of the new landmarks to track, or a different approach is needed.

If the map is not central, an approach based on differential landmark processing could be taken. Instead of tracking the same feature over a sequence of scans, features are only matched between two scans to compute the relative movement between the sweeps.

### Relative Movement Estimation

As described in Section 3.4, all features from the entire scan should not be used to estimate the relative movement at the same time. If the vessel is accelerating or turning the scan will be skewed causing estimation errors. The idea is therefore to use subsets of features, measured over a short time interval, to compute the relative changes in course $\Delta\psi_t$ and position $\Delta X_t$ and $\Delta Y_t$ between two scans. The scans are simply divided into multiple slices where each segment covers a time interval $\tau_t$. The relative movement and course estimates are therefore calculated multiple times per scan pair.

The relative change in position and course can be described as a relationship between the landmark positions measured in the local coordinate frame at time $t$ and $t - T$. These landmark positions are related as

$$\begin{pmatrix} \mathsf{y}_{x,t-T}^i \\ \mathsf{y}_{y,t-T}^i \end{pmatrix} = \begin{pmatrix} \cos(\Delta\psi_t) & -\sin(\Delta\psi_t) \\ \sin(\Delta\psi_t) & \cos(\Delta\psi_t) \end{pmatrix} \begin{pmatrix} \mathsf{y}_{x,t}^i \\ \mathsf{y}_{y,t}^i \end{pmatrix} + \begin{pmatrix} \Delta X_t \\ \Delta Y_t \end{pmatrix} \tag{9}$$

where $\mathsf{y}_{x,t}^i$ is the measured $x-$coordinate of landmark $i$ at time instant $t$ in the local coordinate system. $\mathsf{y}_{x,t-T}^i$ is the measured $x-$coordinate in the previous scan.

If (9) was used to estimate the changes in course and position between two scans using each segment independently, quite large course changes could be experienced. Since each scan pair is used multiple times because it is divided into segments, practically the same course and position change would be calculated over and over again. For example, the change in course registered between the

scans will be similar for two adjacent segments. The only truly new information in the next segment are the changes experienced over that specific segment, not the changes experienced over the rest of the full scan because that has already been studied. To avoid calculating the same course change multiple times, the changes in course and position can be calculated recursively as

$$\Delta\psi_t = \Delta\psi_{t-\tau} + \delta\psi_t \tag{10a}$$
$$\Delta X_t = \Delta X_{t-\tau} + \delta X_t \tag{10b}$$
$$\Delta Y_t = \Delta Y_{t-\tau} + \delta Y_t. \tag{10c}$$

The change in course is subsequently divided into two parts: the estimated change in course $\Delta\psi_{t-\tau}$ using the previous segment, which is known, and a small change in course $\delta\psi_t$ experienced during the segment, which is unknown. Even though the vessel used for data acquisition is very maneuverable, $\delta\psi_t$ can be assumed small.

The sine and cosines of (9) can now be rewritten using (10a) and simplified using the small angle approximation $\cos(\delta\psi_t) \approx 1$ and $\sin(\delta\psi_t) \approx \delta\psi_t$

$$\begin{aligned}
\cos(\Delta\psi) &= \cos(\Delta\psi_{t-\tau} + \delta\psi_t) \\
&= \underbrace{\cos(\Delta\psi_{t-\tau})}_{c_\Delta}\underbrace{\cos(\delta\psi_t)}_{\approx 1} - \underbrace{\sin(\Delta\psi_{t-\tau})}_{s_\Delta}\underbrace{\sin(\delta\psi_t)}_{\approx\delta\psi_t} \\
&\approx c_\Delta - s_\Delta\delta\psi_t \qquad (11) \\
\sin(\Delta\psi) &= \sin(\Delta\psi_{t-\tau} + \delta\psi_t) \\
&= \underbrace{\sin(\Delta\psi_{t-\tau})}_{s_\Delta}\underbrace{\cos(\delta\psi_t)}_{\approx 1} + \underbrace{\cos(\Delta\psi_{t-\tau})}_{c_\Delta}\underbrace{\sin(\delta\psi_t)}_{\approx\delta\psi_t} \\
&\approx s_\Delta + c_\Delta\delta\psi_t \qquad (12)
\end{aligned}$$

where $c_\Delta$ and $s_\Delta$ are known.

Dividing the change in position into two parts as in (10b) and (10c) does not change the equation system (9) in practice, why $\Delta X_t$ and $\Delta Y_t$ are left as before. The equation system becomes

$$\begin{aligned}
\begin{pmatrix} y^i_{x,t-T} \\ y^i_{y,t-T} \end{pmatrix} &= \begin{pmatrix} c_\Delta - s_{\Delta\psi_t}\delta\psi_t & -s_\Delta - c_\Delta\delta\psi_t \\ s_\Delta + c_\Delta\delta\psi_t & c_\Delta - s_\Delta\delta\psi_t \end{pmatrix}\begin{pmatrix} y^i_{x,t} \\ y^i_{y,t} \end{pmatrix} \\
&\quad + \begin{pmatrix} \Delta X_t \\ \Delta Y_t \end{pmatrix} \quad \Leftrightarrow \\
\begin{pmatrix} y^i_{x,t-T} \\ y^i_{y,t-T} \end{pmatrix} &= \left( \begin{pmatrix} c_\Delta & -s_\Delta \\ s_\Delta & c_\Delta \end{pmatrix} + \delta\psi_t \begin{pmatrix} -s_\Delta & -c_\Delta \\ c_\Delta & -s_\Delta \end{pmatrix} \right)\begin{pmatrix} y^i_{x,t} \\ y^i_{y,t} \end{pmatrix} \\
&\quad + \begin{pmatrix} \Delta X_t \\ \Delta Y_t \end{pmatrix} \qquad (13)
\end{aligned}$$

which can be rewritten as

$$\begin{pmatrix} y^i_{x,t-T} - c_\Delta y^i_{x,t} + s_\Delta y^i_{y,t} \\ y^i_{y,t-T} - s_\Delta y^i_{x,t} - c_\Delta y^i_{y,t} \end{pmatrix} = \begin{pmatrix} -s_\Delta y^i_{x,t} - c_\Delta y^i_{y,t} & 1 & 0 \\ c_\Delta y^i_{x,t} - s_\Delta y^i_{y,t} & 0 & 1 \end{pmatrix} \begin{pmatrix} \delta\psi_t \\ \Delta X_t \\ \Delta Y_t \end{pmatrix} \tag{14}$$

The equation system is now approximately linear and by stacking multiple landmarks in one equation system

$$\begin{pmatrix} y^i_{x,t-T} - c_\Delta y^i_{x,t} + s_\Delta y^i_{y,t} \\ y^i_{y,t-T} - s_\Delta y^i_{x,t} - c_\Delta y^i_{y,t} \\ y^j_{x,t-T} - c_\Delta y^j_{x,t} + s_\Delta y^j_{y,t} \\ y^j_{y,t-T} - s_\Delta y^j_{x,t} - c_\Delta y^j_{y,t} \\ \vdots \end{pmatrix} = \begin{pmatrix} -s_\Delta y^i_{x,t} - c_\Delta y^i_{y,t} & 1 & 0 \\ c_\Delta y^i_{x,t} - s_\Delta y^i_{y,t} & 0 & 1 \\ -s_\Delta y^j_{x,t} - c_\Delta y^j_{y,t} & 1 & 0 \\ c_\Delta y^j_{x,t} - s_\Delta y^j_{y,t} & 0 & 1 \\ \vdots & & \end{pmatrix} \begin{pmatrix} \delta\psi_t \\ \Delta X_t \\ \Delta Y_t \end{pmatrix} \tag{15}$$

an overdetermined system is acquired and $\delta\psi_t$, $\Delta X_t$ and $\Delta Y_t$ can be determined using a least squares solver.

This estimated change in position and course can in turn be used to calculate a velocity and an angular velocity measurement as

$$\bar{y}_t = \begin{pmatrix} \bar{v}_t \\ \bar{\omega}_t \end{pmatrix} = \begin{pmatrix} \frac{\sqrt{(\Delta X_t)^2 + (\Delta Y_t)^2}}{T} \\ \frac{\delta\psi_t}{\tau_t} \end{pmatrix} + \begin{pmatrix} \bar{e}_{v,t} \\ \bar{e}_{\omega,t} \end{pmatrix}. \tag{16}$$

The measurement noises $\bar{e}_v$ and $\bar{e}_\omega$ are assumed to be independent Gaussian noises. This transformation that provides direct measurements of speed and course change, gives what is usually referred to as odometry.

Although this approach simplifies the implementation a lot, it comes with certain drawbacks. First, the landmarks correctly associated between two images are used only pair-wise, and this is sub-optimal since one loses both the averaging effects that occur when the same landmark is detected many times and also the correlation structure between landmarks. Second, assuming no cross-correlation between $\bar{e}_v$ and $\bar{e}_\omega$ is a simplification since $\bar{v}_t$ and $\bar{\omega}_t$ are based on $\Delta X_t$, $\Delta Y_t$ and $\delta\psi_t$ which are not calculated independently. Therefore the measurements $\bar{v}_t$ and $\bar{\omega}_t$ are actually dependent making the noise independence assumption incorrect. And third, in order to estimate the relative movements, the time interval used to detect the landmarks must be informative enough to calculate $\delta\psi_t$, $\Delta X_t$ and $\Delta Y_t$, but not long enough to allow a significant scan skewness to appear. This tradeoff is vessel specific and must be balanced. By ensuring that the vessel can not be expected to turn more than for example 10 degrees during each time interval, the small angle approximation holds.

### ESDF

The simplified odometric model above can still be used for mapping if a trajectory based filtering algorithm is used. One such framework is known in the SLAM literature as the Exactly Sparse Delayed-state Filter (ESDF) Eustice et al. (2006).

It has a state vector that consists of augmented vehicle states as

$$z_{1:t} = \begin{pmatrix} z_{1:t-1} \\ z_t \end{pmatrix}, \tag{17}$$

where the state $z_t$ is given in (4) and $z_{1:t-1}$ are all previous poses. If no loop closures are detected, then the ESDF is simply stacking all pose estimates, but once a loop closure is detected and the relative pose between the two time instances is calculated, the ESDF allows for the whole trajectory to be updated using this new information.

Once the trajectory has been estimated, all RADAR scans can be mapped to world coordinates. By overlaying the scans on the estimated trajectory, a RADAR map is obtained. Each pixel now describes how many RADAR detections that have occurred in that coordinate.

## 4 SIFT Performance on RADAR Images

SIFT is used to extract visual island features from the RADAR images. Figure 4 shows the features that are extracted from the upper right quadrant of a RADAR scan example. Two types of features are detected; island features and vessel related features. The latter originate from RADAR disturbances caused by the vessel and the waves and are visible in the bottom left corner of Figure 4. Unfortunately, this section of the image cannot just be removed since the vessel commonly travels very close to land making island features in the area sometimes crucial for navigation.

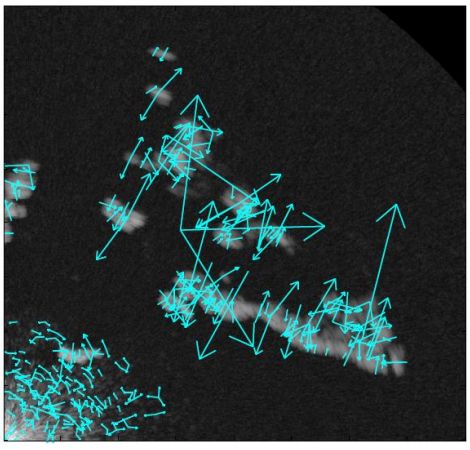

***Figure 4:*** *RADAR image with extracted features marked with arrows. This is a typical example of how a RADAR image from which SIFT extracts a large number of features, may look.*

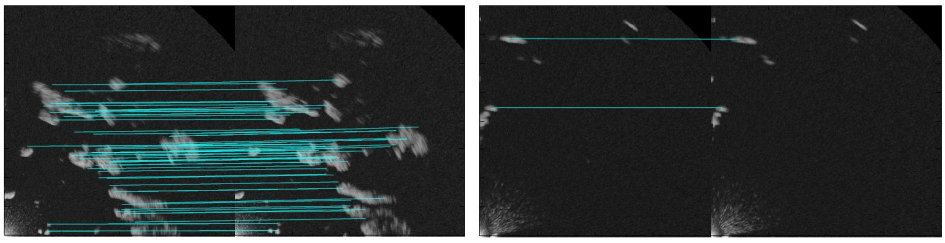

***(a)*** *Plenty of islands makes it easy to find good matching features.*

***(b)*** *Few islands result in few features to match.*

***Figure 5:*** *Two pairs of consecutive RADAR images with detected matches marked with lines.*

The total number of features detected is of course depending on the number of islands in the area, but also on where these islands are situated. A large island close to the vessel will block a large section of the RADAR scan, resulting in few features. In these experiments, an average of 650 features were extracted per full RADAR scan.

## 4.1 Matching for Movement Estimation

The SIFT features are matched to estimate the relative movement of the vessel between multiple consecutive scans. Figures 5a and 5b show examples of how well these features actually match. In Figure 5a, a dense island distribution results in a lot of matches which provide a good movement estimation. In Figure 5b there are very few islands making it difficult to estimate the movements accurately.

There are two situations that can cause few matches. One is when there are few islands and the other is when a large island is very close to the vessel, blocking the view of all the other islands. When the vessel passes close to an island at high speed, the RADAR scans can differ quite significantly between two revolutions. This results not only in few features to match but also in features that can be significantly more difficult to match causing the relative movement estimates to degrade. On average though, about 100 features are matched in each full scan.

## 4.2 Loop Closure Matching

Radar features can also be used to detect loop closures which would enable the filter to update the entire vessel trajectory. The rotation invariance of the SIFT features makes RADAR scans acquired from different headings straightforward to match. Quite a large difference in position is also manageable due to the range of the RADAR sensor. This requires of course that no island is blocking or disturbing the view. Figure 6a shows example locations $a$, $b$ and $c$ that were used to investigate the matching performance of the visual features.

In area $a$, Figure 6b, and $b$, Figure 6c, the features are easy to match despite the

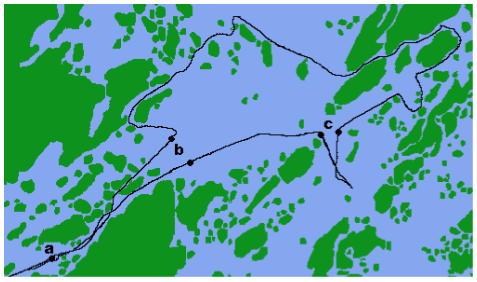


***(a)*** *Example locations used to examine the loop closure detection performance.*

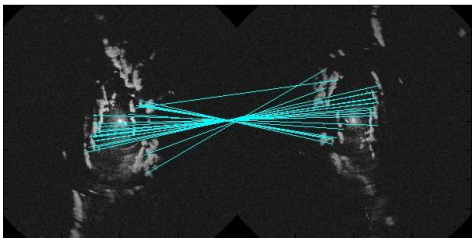

***(b)*** *Multiple matched features from a loop closure in area a in Figure 6a. A 180° difference in course is easily handled by the SIFT features.*

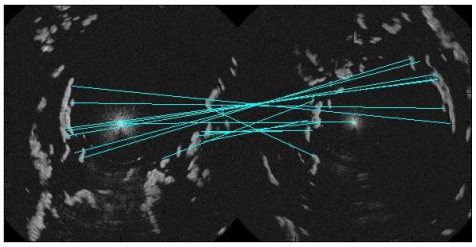

***(c)*** *Location b has both a 180° course difference and a translative change but a couple of features are still matched.*

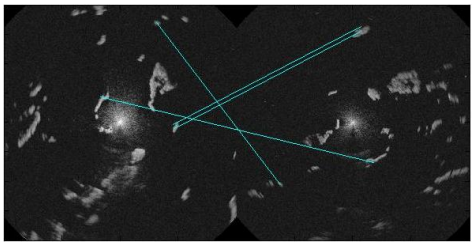

***(d)*** *2 correct and 2 false matches in area c makes it unsuitable for loop closure updates.*

***Figure 6:*** *Examples of loop closure experiments using visual SIFT features.*

rather long translational difference over open water in *b*. In both cases a 180° difference in course is easily overcome by the visual features. This shows the strength of both the RADAR sensor and of the visual features. The long range of the sensor makes loop closures over a wide passage of open water possible. These scans would be used in a full scale SLAM experiment to update the trajectory.

In area *c*, Figure 6d, only two features are matched correctly and there are also two false positives. If the scans are compared ocularly it is quite challenging to find islands that are clearly matching, mostly due to blurring and to blocking islands. It is also noticeable that the RADAR reflections from the islands differ due to differences in RADAR positions which of course alters the SIFT features. The poor matching result is therefore natural in this case.

## 4.3 Feature Preprocessing

Two problems remain that have not been addressed. First is the problem of false feature matches. In Figure 6b a false feature match is clearly visible and it would introduce estimation errors if not handled properly. An initial approach would be to use an algorithm like RANSAC Fischler and Bolles (1981) to remove matches that are inconsistent with all other matches. One could also use the filtering

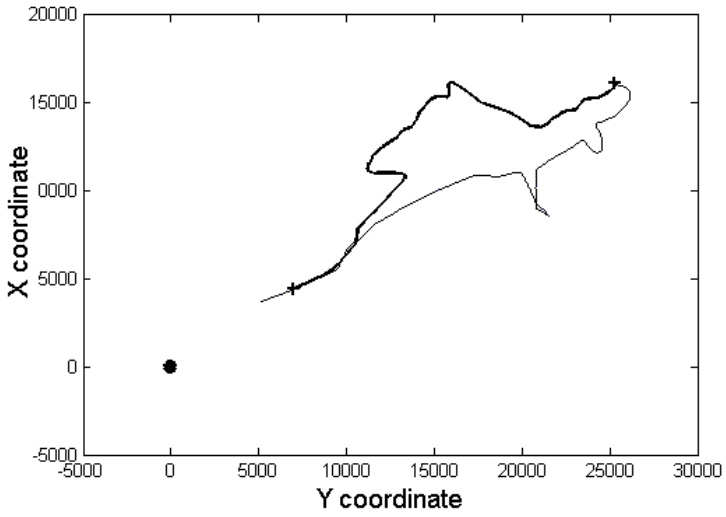


***Figure 7:*** *The whole trajectory where the studied interval is marked by bold line.*

framework to produce an estimate of the probable feature position and perform a significance test on the features based on for example Mahalanobis distance. Only the features that pass this test would be allowed to update the filter estimates.

The other problem is accidental vessel matching. In heavily trafficked areas like ports, other vessels will be detected on the radar scans. If a moving vessel is deemed stationary and is used to update the filter, errors will be introduced. Two approaches could be taken to handle this problem. Again, a significance test could be used to rule out features from fast moving vessels. Alternatively, a target tracking approach could be used. By tracking the features over multiple scans, the vessels can be detected and ruled out based on their position inconsistency compared to the stationary features. Describing such a system is though beyond the scope of this paper. The joint SLAM and target tracking problem has previously been studied in Wang et al. (2007).

## 5 Experimental Results

The experimental results in this section come from the master thesis by Henrik Svensson Svensson (2009). The implemented framework is the one described in Sections 3.5 and 3.5.

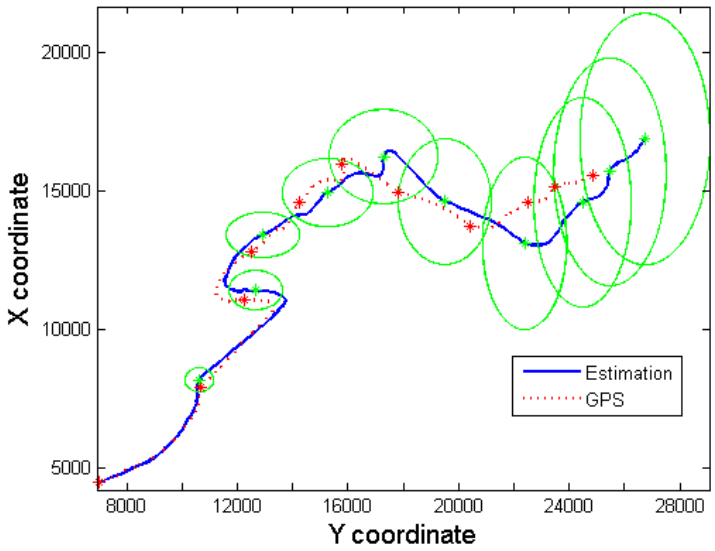


***Figure 8:*** *ESDF trajectory estimate with covariance ellipses compared to GPSmeasurements.*

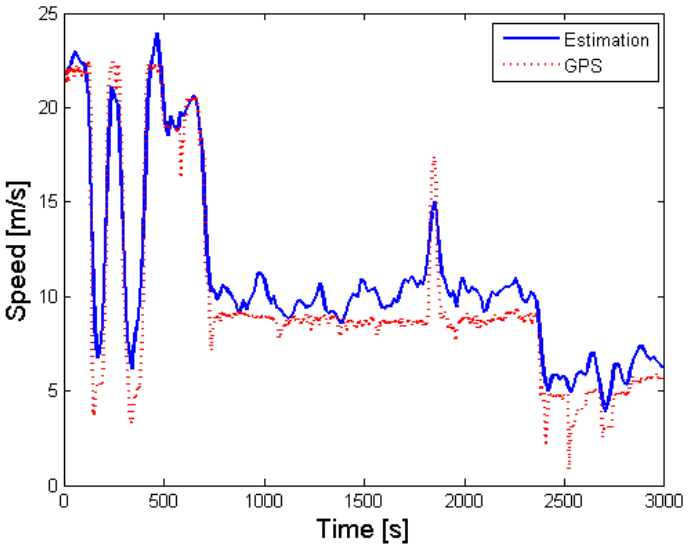


***Figure 9:*** *Velocity estimate compared to GPSmeasurements. A slight positive bias is present, probably due to the simplifications mentioned in Section 3.5.*

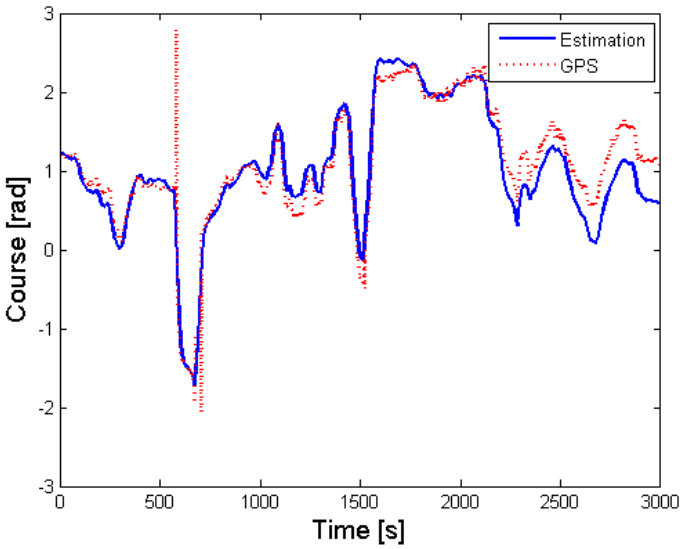


***Figure 10:*** *Course estimate compared to GPSmeasurements. The course estimate is the sum of all estimated changes in course, causing the error to grow in time.*

## 5.1 Results

The trajectory used in the SLAM experiment is shown in bold in Figure 7. The track is about 3000 seconds long (50 minutes) and covers roughly 32 km. The entire round trip was unfortunately never used in one single experiment since it was constituted of multiple data sets.

The estimated trajectory with covariance is compared to the GPS data in Figure 8. The first two quarters of the trajectory consists of an island rich environment, see Figure 6a, resulting in a very good estimate. The third quarter covers an area with fewer islands causing the performance to degrade. This results in an initial misalignment of the final quarter which makes the estimated trajectory of this segment seem worse than it actually is.

Both velocity and course estimates, Figures 9 and 10, are quite good when compared to GPS data. There is though a positive bias on the velocity estimate, probably due to the simplifications mentioned in Section 3.5. The course error grows in time since the estimate is the sum of a long sequence of estimated changes in course, see (16), and there are no course measurements available. The final course error is about 30 degrees.

### 5.2 Map Estimate

A RADAR map of the area was generated by overlaying the RADAR scans on the estimated trajectory. Figure 11 shows the estimated RADAR map that should be compared to the map created using the GPS trajectory, Figure 12. They are very similar although small errors in the estimated trajectory are visible in Figure 11 as blurriness. Some islands appear a bit larger in the estimated map because of this. Overall the map estimate is good.

The estimated map should also be compared to the satellite photo of the area with the true trajectory marked in white as shown in Figure 13. When compared, many islands in the estimated map are easily identified. This shows that the rather simple approach of using visual features on RADAR images can provide good mapping results.

## 6 Conclusions

We have presented a new approach to robust navigation for surface vessels based on RADAR measurements only. No infrastructure or maps are needed. The basic idea is to treat the RADAR scans as images and apply the SIFT algorithm for tracking landmarks of opportunity. We presented two related frameworks, one based on the SLAM idea where the trajectory is estimated jointly with a map, and the other one based on odometry. We have evaluated the SIFT performance and the odometric framework on data from a high-speed patrol boat, and obtained a very accurate trajectory and map.

An interesting application of this work would be to apply this method to underwater vessels equipped with synthetic aperture sonar as the imagery sensor, since there are very few low-cost solutions for underwater navigation.

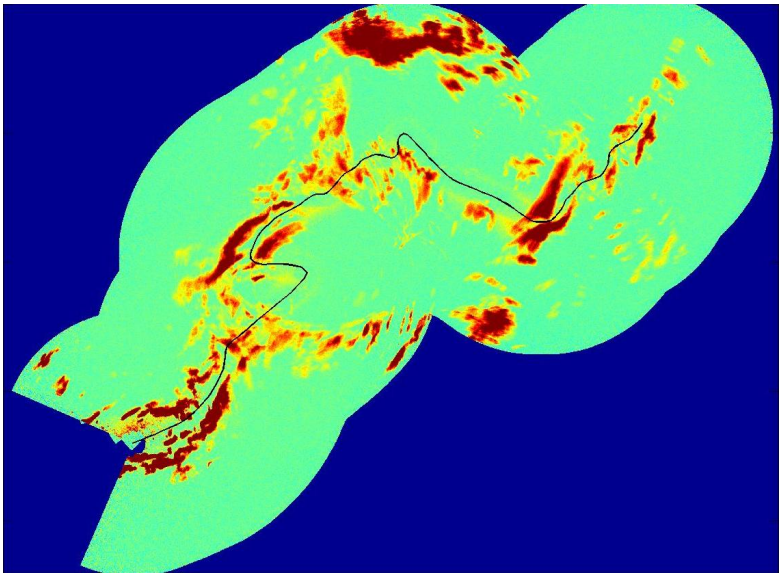

***Figure 11:*** *The estimated map created by overlaying the RADAR scans on the estimated trajectory. The estimated trajectory is marked with a black line.*

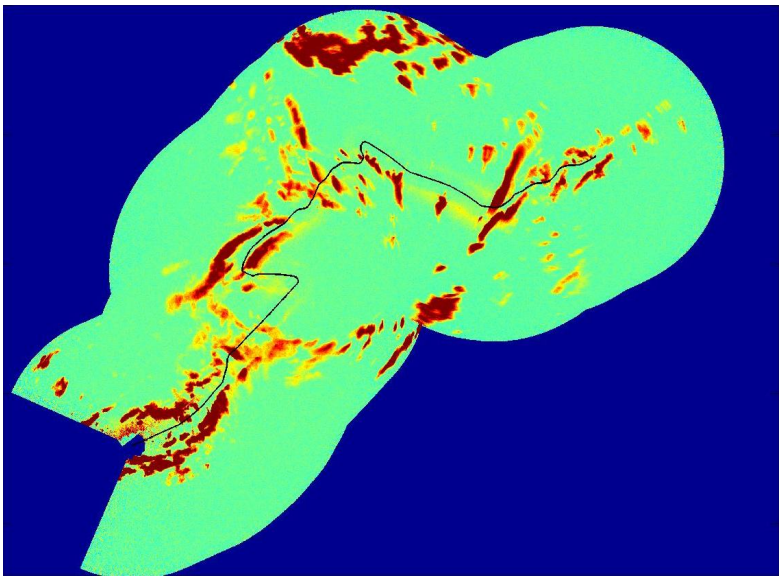

***Figure 12:*** *A RADAR map created using the GPS measurements of the trajectory. The true trajectory is marked with a black line.*

***Figure 13:*** *A satellite photo of the area. The true trajectory is marked with a white line. Many islands in Figure 11 are easily identified in this photo.*

## Bibliography

GNSS vulnerability and mitigation measures, (rev. 6). Technical report, European Maritime Radionavigation Forum, 2001.

T. Bailey and H. Durrant-Whyte. Simultaneous localization and mapping (SLAM): part II. *Robotics & Automation Magazine, IEEE*, 13(3):108 – 117, 2006.

H. Bay, T. Tuytelaars, and L Van Gool. Surf: Speeded up robust features. In *9th European Conference on Computer Vision*, Graz Austria, May 2006.

S. M. Bolen and V. Chandrasekar. Methodology for aligning and comparing spaceborne radar and ground-based radar observations. *J. Atmos. Oceanic Technol.*, 20:647 – 659, 2003.

M. C. Bosse and R. Zlot. Map matching and data association for large-scale two-dimensional laser scan-based SLAM. *International Journal of Robotics Research*, 27(6):667–691, 2008.

M. Bryson and S. Sukkarieh. Bearing-only SLAM for an airborne vehicle. In *Proceedings of the Australasian Conf. Robot. Autom. (ACRA)*, 2005.

J. Callmer, K. Granström, J. Nieto, and F. Ramos. Tree of words for visual loop closure detection in urban SLAM. In *Proceedings of the 2008 Australasian Conference on Robotics and Automation (ACRA)*, 2008.

J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011.

P. Carlbom. Radar map matching. Technical Report, 2005.

M. Chandran and P. Newman. Motion estimation from map quality with millimeter wave radar. In *Proceedings of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2006.

P. Checchin, F. Gérossier, C. Blanc, R. Chapuis, and L. Trassoudaine. Radar scan matching SLAM using the Fourier-Mellin transform. In *IEEE Intl. Conf. on Field and Service Robotics*, 2009.

Q. Chen, M. Defrise, and F. Deconinck. Symmetric phase-only matched filtering of fourier-mellin transforms for image registration and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(12):1156–1168, 1994.

Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.

S. Clark and G. Dissanayake. Simultaneous localisation and map building using millimetre wave radar to extract natural features. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1999.

S. Clark and H. Durrant-Whyte. Autonomous land vehicle navigation using millimeter wave radar. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1998.

M. Cummins and P. Newman. Probabilistic appearance based navigation and loop closing. In *IEEE Intl Conf. on Robotics and Automation*, 2007.

M. Dalin and S. Mahl. Radar map matching. Master's thesis, Linköping University, Sweden, 2007.

H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping (SLAM): part I. *Robotics & Automation Magazine, IEEE*, 13(2):99 – 110, 2006.

H. Essen, G. Luedtke, P. Warok, W. Koch, K. Wild, and M. Schikora. Millimeter wave radar network for foreign object detection. In *2nd International Workshop on Cognitive Information Processing (CIP)*, 2010.

R. Eustice, H. Singh, J. Leonard, M. Walter, and R. Ballard. Visually navigating the RMS titanic with SLAM information filters. In *Proceedings of Robotics: Science and Systems*, 2005.

R.M. Eustice, H. Singh, and J.J Leonard. Exactly sparse delayed-state filters for view-based SLAM. *IEEE Transactions on Robotics*, 2006.

Lu Feng and E.E. Milios. Robot pose estimation in unknown environments by matching 2D range scans. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*, 1994.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24:381–395, June 1981. ISSN 0001-0782.

K. Granström, J. Callmer, F. Ramos, and J. Nieto. Learning to detect loop closure from range data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.

A. Grant, P. Williams, N. Ward, and S. Basker. GPS jamming and the impact on maritime navigation. *Journal of Navigation*, 62(2):173–187, 2009.

F. Gustafsson. *Statistical Sensor Fusion*. Studentlitteratur, 2010.

P. Jensfelt, D. Kragic, J. Folkesson, and M. Bjorkman. A framework for vision based bearing only 3-D SLAM. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2006.

R. Karlsson and F. Gustafsson. Bayesian surface and underwater navigation. *IEEE Transactions on Signal Processing*, 54(11):4204–4213, November 2006.

F. Li, G. Zhang, and J. Yan. Coregistration based on SIFT algorithm for synthetic aperture radar interferometry. In *Proceedings of ISPRS Congress*, 2008.

D.G. Lowe. Object recognition from local scale-invariant features. In *ICCV '99:*

*Proceedings of the International Conference on Computer Vision-Volume 2*, 1999.

Feng Lu and E.E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 1997.

C. Lundquist. Automotive sensor fusion for situation awareness. Linköping University, Sweden, 2009. Licentiate Thesis No 1422.

I. Mahon, S.B. Williams, O. Pizarro, and M. Johnson-Roberson. Efficient view-based SLAM using visual loop closures. *IEEE Transactions on Robotics*, 24(5): 1002–1014, 2008.

P. Newman and K. Ho. SLAM-loop closing with visually salient features. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

F. Ramos, D. Fox, and H.F. Durrant-Whyte. CRF-matching: Conditional random fields for feature based scan matching. In *Proceedings of Robotics: Science and Systems*, 2007a.

F. Ramos, J. Nieto, and H.F. Durrant-Whyte. Recognising and modelling landmarks to close loops in outdoor SLAM. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2007b.

R. Rouveure, M.O. Monod, and P. Faure. High resolution mapping of the environment with a ground-based radar imager. In *Proceedings of the International Radar Conference*, 2009.

M. Schikora and B. Romba. A framework for multiple radar and multiple 2D/3D camera fusion. In *4th German Workshop Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, 2009.

S. Se, D. Lowe, and J. Little. Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. *The International Journal of Robotics Research*, 21:735–758, 2002.

H. Svensson. Simultaneous Localization And Mapping in a Marine Environment using Radar Images. Master's thesis, Linköping University, Sweden, 2009.

S. H. Tsang, P. S. Hall, E. G. Hoare, and N. J. Clarke. Advance path measurement for automotive radar applications. *IEEE Trans. on Intelligent Transportation Systems*, 7(3):273 – 281, 2006.

J. Volpe. Vulnerability assessment of the transportation infrastructure relying on global positioning system final report. Technical report, National Transportation Systems Center, 2001.

Chieh-Chih Wang, Charles Thorpe, Sebastian Thrun, Martial Hebert, and Hugh Durrant-Whyte. Simultaneous localization, mapping and moving object tracking. *The International Journal of Robotics Research*, 26(9):889–916, 2007. doi: 10.1177/0278364907081229.

# Paper E

# Silent Localization of Underwater Sensors using Magnetometers

*Authors:* Jonas Callmer, Martin Skoglund and Fredrik Gustafsson

*Edited version of the paper:*

J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010a.

# Silent Localization of Underwater Sensors using Magnetometers

Jonas Callmer, Martin Skoglund and Fredrik Gustafsson

Dept. of Electrical Engineering,
Linköping University,
SE–581 83 Linköping, Sweden
{callmer, ms, fredrik}@isy.liu.se

**Abstract**

Sensor localization is a central problem for sensor networks. If the sensor positions are uncertain, the target tracking ability of the sensor network is reduced. Sensor localization in underwater environments is traditionally addressed using acoustic range measurements involving known anchor or surface nodes. We explore the usage of triaxial magnetometers and a friendly vessel with known magnetic dipole to silently localize the sensors. The ferromagnetic field created by the dipole is measured by the magnetometers and is used to localize the sensors. The trajectory of the vessel and the sensor positions are estimated simultaneously using an Extended Kalman Filter (EKF). Simulations show that the sensors can be accurately positioned using magnetometers.

## 1 Introduction

Today, surveillance of ports and other maritime environments is getting increasingly important for naval and customs services. Surface vessels are rather easy to detect and track, unlike submarines and other underwater vessels which pose new threats such as terrorism and smuggling. To detect these vessels, an advanced underwater sensor network is necessary. Such sensors can measure fluctuations in for example magnetic and electric fields, pressure changes and acoustics.

Deploying an underwater sensor in its predetermined position can be difficult due to currents, surge and the lack of a Global Navigation Satellite System (GNSS) functioning underwater. Sometimes the sensors must be deployed fast, resulting in very uncertain sensor positions. These positions must then be estimated in order to enable the network to accurately track an alien vessel.

Lately, many solutions to the underwater sensor localization problem have been suggested. They can be broadly divided into two major categories: range-based

and range-free. In general, range-based schemes provide more accurate positioning than range-free schemes.

Range-based schemes use information about the range or angle between sensors. The problem is thereafter formulated as a multilateral problem. Common methods to measure range or angle include Time of Arrival (ToA), Time Difference of Arrival (TDoA), Angle of Arrival (AoA) or Received Signal Strength Indicator (RSSI). These methods usually require active pinging but silent methods based on TDoA have been suggested Cheng et al. (May. 2008). The 3D positioning problem can be transformed into a 2D problem by the use of depth sensors Cheng et al. (2008). The range positioning scheme is often aided by surface nodes, anchor nodes, mobile beacons or autonomous underwater vehicles Zhou et al. (2007); Erol et al. (2008, 2007). Joint sensor localization and time synchronization was performed in Tian et al. (2007).

Range-free schemes generally provide a coarse estimate of a node's location and their main advantage lie in their simplicity. Examples of range-free schemes are Density-aware Hop-count Localization (DHL) Wong et al. (2005) and Area Localization Scheme (ALS) Chandrasekhar and Seah (2006). A more thorough description of underwater sensor localization solutions, can be found in the surveys Akyildiz et al. (2005); Chandrasekhar et al. (2006).

In this paper we propose a method to silently localize underwater sensors equipped with triaxial magnetometers using a friendly vessel with known static magnetization characteristics. (For methods to estimate the magnetic characteristics, see Nelson and Richards (2007).) Each sensor is further equipped with a pressure sensor and an accelerometer used for depth estimation and sensor orientation estimation, respectively. To enable global positioning of the sensors, the vessel or one sensor must be positioned globally. To the best of the authors knowledge this is the first time magnetic dipole tracking is used for sensor localization.

For target tracking in shallow waters, magnetometers are often a more useful sensor than acoustics, since sound scatters significantly in these environments Birsan (2006). Birsan has explored the use of magnetometers and the magnetic dipole of a vessel for target tracking, Birsan (2004, 2005). Two sensors with known positions were used to track a vessel while simultaneously estimating the unknown magnetic dipole of the vessel. Tracking and estimation were performed using an unscented Kalman filter Birsan (2004) and an unscented particle filter Birsan (2005). Dalberg et al. (2006) fused electromagnetic and acoustic data to track surface vessels using underwater sensors.

Several studies of the electromagnetic characteristics of the maritime environment have stated that the permeability of the seabed differs considerably from the permeability of air or water. The environment should therefore be modeled as a horizontally stratified model with site specific permeability and layer thickness for each segment Dalberg et al. (2006); Daya et al. (2005); Birsan (2006). This has not been included in our simulation study but should be considered in field experiments.

In the past 10-15 years quite a lot of effort has been put into reducing the static magnetic signature of naval vessels by active signature canceling. This has increased the importance of other sources of magnetic fields such as Corrosion Related Magnetism (CRM) Daya et al. (2005); Lundin (2003). CRM is generated by currents in the hull, normally induced by corrosion or the propeller. It is therefore very difficult to estimate and subsequently difficult to cancel. This makes CRM important in naval target tracking but not so much in sensor localization. In our study, a friendly vessel used for sensor localization can turn off its active signature canceling, resulting in a magnetic field from the dipole which is considerably stronger than the field induced by CRM. We have therefore neglected CRM.

An underwater sensor network used for real time surveillance must be silent. Neither sensor localization, surveillance or data transfer can be allowed to expose the sensor network. Silent communication rules out the use of acoustic modems which are the common mean of wireless underwater data transfer Akyildiz et al. (2005). We therefore assume that the sensors are connected by wire. As a consequence, common problems in underwater sensor networks such as time synchronization, limited bandwidth and limited energy resources will be neglected.

The sensor localization problem is basically reversed Simultaneous Localization and Mapping (SLAM). In common SLAM Durrant-Whyte and Bailey (2006); Bailey and Durrant-Whyte (2006), landmarks in the environment are tracked with on-board sensors. The positions of these landmarks and the vehicle trajectory are estimated simultaneously in a filter. In sensor localization the sensors are observing the vessel from the "landmarks" position. The problem has the same form as SLAM but with a known number of landmarks and known data association.

The paper outline is as follows: Section 2 describes the system, measurement models and state estimation. Simulations of the performance of the positioning scheme, its sensitivity to different errors and the importance of the appearance of the trajectory are studied in Section 3. The paper ends with conclusions.

# 2 Methodology

This section describes the nonlinear state estimation problem here solved with EKF-SLAM, how the vessel dynamics and sensors are modeled and how different performance measures are computed. All vectors are expressed in a world fixed coordinate system unless otherwise stated.

## 2.1 System Description

The sensor positioning system is assumed to have the following process and measurement model

$$\boldsymbol{x}_{k+1} = f(\boldsymbol{x}_k) + \boldsymbol{w}_k \tag{1}$$

$$\boldsymbol{y}_k = h(\boldsymbol{x}_k, \boldsymbol{u}_k, \boldsymbol{e}_k^u) + \boldsymbol{e}_k \tag{2}$$

where $f(\cdot)$ is a nonlinear state transition function, $h(\cdot)$ is a nonlinear measurement function, $\boldsymbol{x}_k$ the state vector, $\boldsymbol{u}_k$ the inputs, $\boldsymbol{w}_k$ the process noise, $\boldsymbol{e}_k^u$ input noise and $\boldsymbol{e}_k$ measurement noise. In SLAM the state vector consists of both the vessel position $\boldsymbol{p}_v = [x, y]^T$ and landmark (sensor) states $\boldsymbol{s}$ stacked, i.e. $\boldsymbol{x} = [\boldsymbol{p}_v^T, \boldsymbol{s}^T]^T$.

### Process Model

The process model describes the vehicle and the sensors dynamics. There are complex vessel models available which include 3D orientation, angular rates, engine speed, rudder angle, waves, hull, etc., see e.g. Fossen and Perez (2004). Since we do not consider any particular vessel or weather condition, a very simple vessel model is used. It is assumed that no substantial movements in the $z$-coordinate, pitch and roll angles of the vessel are made, hence a nonlinear 5 states coordinated turn model is sufficient. The parametrization used is a linearized discretization according to Gustafsson (2001)

$$x_{k+T} = x_k + \frac{2v_k}{\omega_k}\sin(\omega_k T)\cos(h_k + \frac{\omega_k T}{2}) \tag{3a}$$

$$y_{k+T} = y_k + \frac{2v_k}{\omega_k}\sin(\omega_k T)\sin(h_k + \frac{\omega_k T}{2}) \tag{3b}$$

$$v_{k+T} = v_k \tag{3c}$$

$$h_{k+T} = h_k + \omega_k T \tag{3d}$$

$$\omega_{k+T} = \omega_k \tag{3e}$$

where $T$ is the sampling interval and $(x, y)$, $v$, $h$, $\omega$ denote position, speed, heading and angular rate, respectively. Furthermore, it is assumed that the sensors are static and do not move after some time of deployment, hence a process model for the sensors is

$$s_{x_j,k+T} = s_{x_j,k} \tag{4a}$$

$$s_{y_j,k+T} = s_{y_j,k} \qquad j = 1, \ldots, M \tag{4b}$$

where $M$ is the number of sensors, $s_{x_j}$ and $s_{y_j}$ are sensor $j$'s $x$ and $y$ position, respectively.

### Measurement Model

Each sensor contains a pressure sensor which is used as an input, $d_{j,k}$, of the $z$-component. The sensor also contains accelerometers which are used to determine the direction of the earth gravitational field. The magnetometers in the sensor can be used to compute the direction of the earth magnetic inclination if the environment is free from magnetic disturbances such as ships. In most cases the magnetic inclination vector will not be parallel to the gravitational vector (except for the magnetic north and south pole) and the sensor orientation may be readily measured. The sensor orientation is modeled as a static input $\boldsymbol{C}_j$.

In this paper we only consider the ferromagnetic signature due to the iron in the vessel construction. The ferromagnetic signature stems from the large pieces of metal used to construct a vessel. Each piece has its own magnetic dipole and

the sum of these dipoles can roughly be simplified into a single dipole. The magnetic flux density for a dipole diminishes cubically with the distance to the dipole. With vector magnetometers dipole orientation can be estimated. Triaxial measurements of the magnetic flux density from a dipole can be modeled as

$$h(\boldsymbol{x}_k, \boldsymbol{u}_k) = \tag{5}$$

$$\frac{\mu_0}{4\pi|\boldsymbol{r}_{j,k}|^5}(3\langle \boldsymbol{r}_{j,k}, \boldsymbol{m}(h_k)\rangle \boldsymbol{r}_{j,k} - |\boldsymbol{r}_{j,k}|^2 \boldsymbol{m}(h_k)) \tag{6}$$

where $\boldsymbol{m}(h_k) = [m_x \cos(h_k), m_y \sin(h_k), m_z]^T$ is the magnetic dipole of the vessel and $\mu_0$ is the permeability of the medium. $\boldsymbol{r}_{j,k} = \boldsymbol{C}_j[x_k - s_{x_{j,k}}, y_k - s_{y_{j,k}}, 0 - d_{j,k}]^T$ is the vector from each sensor to the vessel where $\boldsymbol{C}_j$ is the static orientation of sensor $j$ in the global coordinate frame and $d_{j,k}$ is the measured depth of the sensor. Note that $d_{j,k}$ and $\boldsymbol{C}_{j,k}$ should be seen as inputs, $\boldsymbol{u}_k = \{d_{j,k}, \boldsymbol{C}_j\}_{j=1}^M$, since these are measured variables but not part of the state vector. The dipole model without coordinate transformations can be found in e.g. Cheng (1989). In the proximity of the vessel, a possibly better model would be a multiple dipole model Lindin (2007) where the measurement is the sum of several dipoles, but this is out of the scope of this paper. A single dipole is a reasonable approximation if the measurements are made at a large distance compared to vessel size Nelson and Richards (2007).

The magnetic dipole used throughout the simulations was $\boldsymbol{m} = [50, -5, 125]^T$ kAm$^2$ (same as in Birsan (2005)). Fig. 1 shows the measured magnetic flux density at sensor 3 in Fig. 2 from a vessel where the dipole has been slightly rotated around the $z$-axis between each simulation. The upper left figure in Fig. 1 was acquired using the magnetic dipole discussed earlier. Clearly the direction of the dipole affects the measured magnetic field. This indicates the importance of using a accurate dipole estimate.

## 2.2 State Estimation

Our approach to the state estimation problem is to use an Extended Kalman Filter (EKF) in the formulation of EKF-SLAM, for details see e.g. Durrant-Whyte and Bailey (2006). There are some characteristics in this system which do not usually exist in the common SLAM problem:

- The landmarks (sensor) are naturally associated to the measurements, i.e. data association is solved.
- The sensors global orientations are known which in turn makes it possible to estimate the orientation of the trajectory.
- The planar motion assumption and the pressure sensor makes it possible to transform sensor positioning into 2D.

A well known problem with SLAM is the ever expanding state space that comes with addition of new landmarks which will eventually make it intractable to compute a solution. In a sensor network the number of sensors (landmarks) will nor-

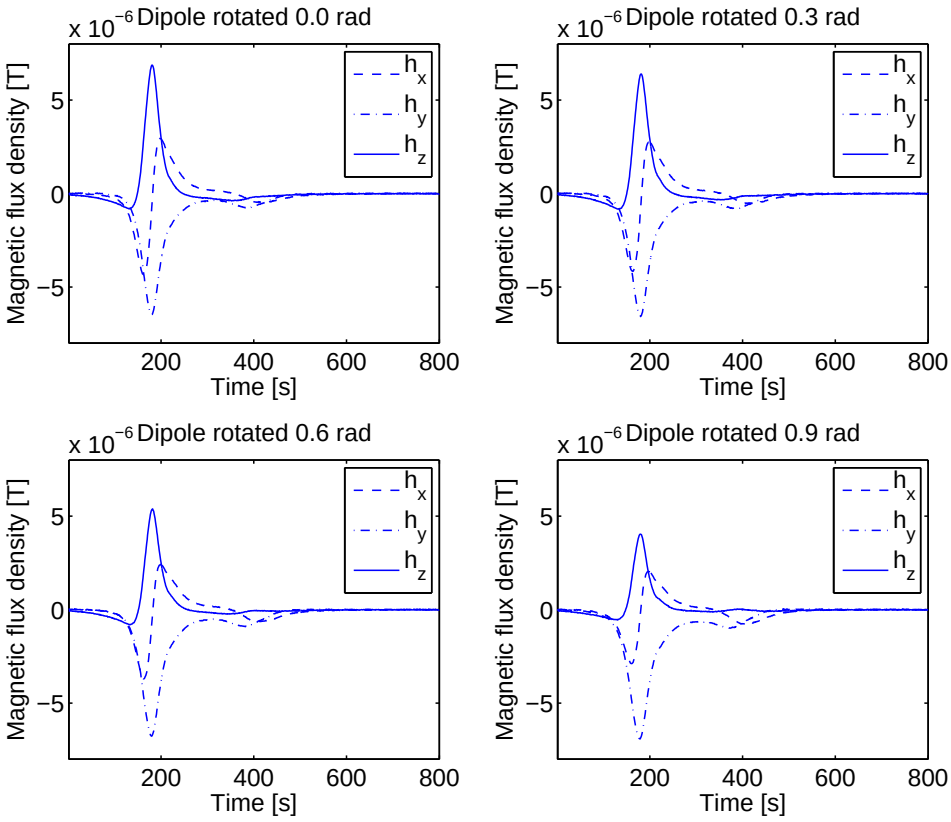


***Figure 1:*** *Measured magnetic flux density at sensor 3 in Fig. 2 for vessels with slightly rotated dipoles.*

mally be known.

Due to the duality of the estimation problem implied in SLAM, i.e. *estimate a map and simultaneously localize the vehicle in the map*, the question of state observability needs to be answered. Previous observability analyzes on the SLAM problem Kim and Sukkarieh (2004); Bryson and Sukkarieh (2008); Andrade-Cetto and Sanfeliu (2005); Lee et al. (2006); Wang and Dissanayake (2008); Perera et al. (2009) has focused on vehicle fixed range and/or bearing sensors, such as laser and camera. Perera et al. (2009); Andrade-Cetto and Sanfeliu (2005) conclude that only one known landmark needs to be observed in 2D SLAM for the global map to be locally weakly observable. In our proposed system the sensors are in the actual landmarks position and their measurements are informative in both range and bearing to the dipole, hence the global map is observable if one sensor position is known. Theoretically this means that the sensor positions and the trajectory can be estimated in a global coordinate frame with a global map position error depending only on the error of the known sensor. If no global position of either sensor or vessel is available the sensors can be positioned locally.

Even if the system is observable there are no guarantees that an EKF will converge since it depends on the linearization error and the initial linearization point. More recent approaches to the SLAM problem Kaess et al. (2007); Dellaert (2005) consider smoothing instead of filtering. These methods can handle linearization errors better since the whole trajectory and map can be re-linearized. Yet, a good initial linearization point is necessary.

## 2.3 Cramer-Rao Lower Bound

Given the trajectory of a vessel, it is interesting to study a lower bound on the covariance of the estimated sensor positions. We have chosen to study the Cramer-Rao lower bound (CRLB) due to its simplistic advantages. CRLB is the inverse of the Fisher Information Matrix (FIM), $I(\boldsymbol{x})$, which in case of Gaussian measurement errors can be calculated as

$$I(\boldsymbol{x}) = H(\boldsymbol{x})^T R(\boldsymbol{x})^{-1} H(\boldsymbol{x}), \tag{7a}$$

$$H(\boldsymbol{x}) = \nabla_{\boldsymbol{x}} h(\boldsymbol{x}) \tag{7b}$$

where $R(\boldsymbol{x})$ is Gaussian measurement noise and $H(\boldsymbol{x})$ denote the gradient of $h(\boldsymbol{x})$ w.r.t. $\boldsymbol{x}$. The CRLB of a sensor position is given by

$$Cov(\boldsymbol{s}) = E\left\{(\boldsymbol{s^0} - \boldsymbol{s})(\boldsymbol{s^0} - \boldsymbol{s})^T\right\} \tag{8a}$$

$$\geq I(\boldsymbol{s})^{-1} \tag{8b}$$

where $\boldsymbol{s^0}$ is true sensor position and $\boldsymbol{s}$ is the corresponding estimate. Since information is additive the FIM of a sensor location for a certain trajectory can be calculated as the sum of the FIMs from all vessel positions along the trajectory. The lower bound of the covariance of the sensor position estimate is then the inverse of the sum of the FIMs. A more extensive study of the fundamentals of CRLB can be found in Kay (1993).

# 3 Simulation Results

The sensor positioning problem can, depending on which sensors are available, be solved in different ways. If no accurate global position of the vessel or a sensor is available during the experiment (GPS is for example easily jammed), the sensors can only be positioned locally. In Section 3.1, magnetometers are used to localize the sensors. If global vessel position is available throughout the experiment, from GNSS or using a radar sensor and a sea chart, it can be used as a measurement of the position of the vessel. This will not only position the sensors globally but also enable a more accurate trajectory estimation. This experimental configuration is simulated in Section 3.2. The parameters used in the simulations are listed in Table 1.

| **Param.** | Covariance SLAM/GNSS | **Param.** | Value |
|---|---|---|---|
| $x_0$ | 10/10 m | $\boldsymbol{m}$ | [50, -5, 125] kAm$^2$ |
| $y_0$ | 10/10 m | $\mu_0$ | $4\pi$ $10^{-7}$ Tm/A |
| $v_0$ | 0/0 m/s | $d_{j,0}$ | {-5, -15} m |
| $h_0$ | 1/1 rad | $T$ | 0.1 s |
| $\omega_0$ | 0/0 rad/s | | |
| $s_{x_j}$ | 400/400 m | | |
| $e_{GNSS}$ | 1 m | | |
| $e_h$ | $10^{-16}$ T | | |

***Table 1:*** *The parameters used in the simulations.*

## 3.1 Magnetometers Only

If there is no reliable global position measurement of the vessel, the trajectory of the vessel must be estimated using the same magnetic fluctuations as are being used to localize the sensors. Simulations show that the sensor network needs to be more dense when no GNSS is available. If there is little or no overlap in which two or more sensors observe the vessel simultaneously, the trajectory estimate, and in the end the sensor position estimates, depend more on the vessel model than observations. Yet, the sensor positions are still coupled through the covariance matrix.

A sensor localization simulation using 7 sensors and a generated trajectory is shown in Fig. 2 and Fig. 3. Fig. 4 shows the Root Mean Square Error (RMSE) of each sensor as it develops over time. Since the initial guesses of sensor positions were generated independently, different sensors have different initial errors. All sensor though have the same initial uncertainty covariance (400 m, see Table 1 ). The initial guesses are meant to represent the prior information of the true sensor positions, acquired during sensor deployment. The limited range of the magnetic fluctuations causes the sensor position estimate to change only when the vessel is sufficiently close. This can be studied in Fig. 4. Sensor 4 in Fig. 2 is too far

away from the vessel for accurate positioning, resulting in a large uncertainty ellipse. From Fig. 2, it is clear that error in trajectory estimates result in errors in estimated sensor positions.

200 Monte Carlo simulations using different trajectories and sensor locations show that this configuration results in a positioning error of 26.3% in average. A sensor failing to retain the true sensor position within two standard deviations was considered incorrectly positioned. In Fig. 2, sensor 7 is incorrectly positioned.

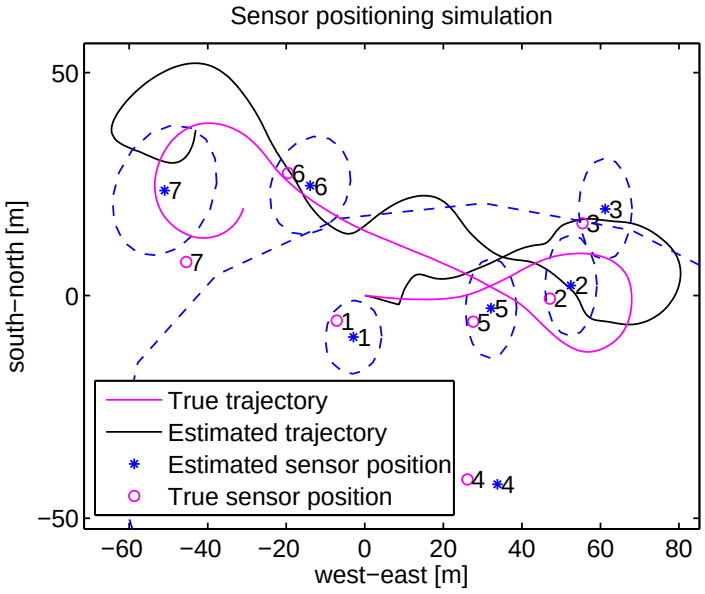


***Figure 2:*** *Estimated sensor positions with 2σ uncertainty and vessel trajectory, for simulations using magnetometers.*

## 3.2 Magnetometers and GNSS

If global position measurements of the vessel are available throughout the trajectory, these measurements are used to improve the trajectory estimate. Each sensor is positioned relative to the trajectory of the vessel and is therefore less dependent of other sensor positions than in Section 3.1. This is quite natural since the cross correlations will not have such great impact on the sensor position estimates when the trajectory is known. Simulation results using the same sensor positions and trajectory used in Section 3.1, are shown in Fig. 5. Fig. 6 shows the RMSE of each sensor throughout the simulation. The global trajectory measurements result in more accurate sensor position estimates and lower uncertainties than using only magnetometers. Sensor 4 is far away from the trajectory resulting in a very uncertain position estimate.

200 Monte Carlo simulations using different trajectories and sensor locations

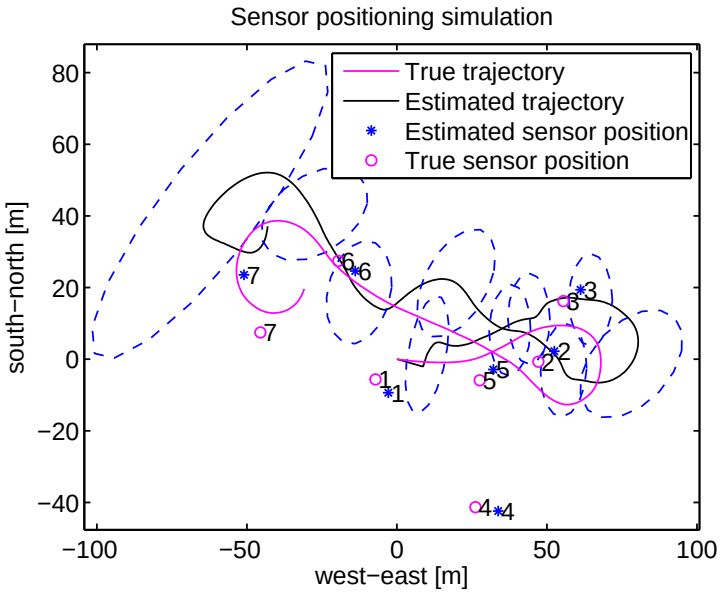


***Figure 3:*** *Estimated vessel trajectory with $2\sigma$ uncertainty and sensor positions, for simulations using magnetometers.*

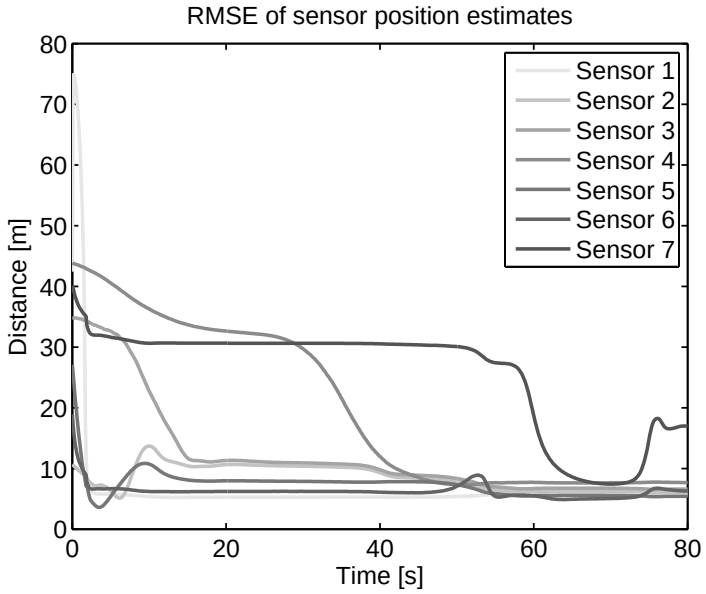


***Figure 4:*** *Root Mean Square Error of estimated sensor positions throughout the simulation using only magnetometers.*

show that using magnetometers and GNSS results in a sensor positioning error of 12.9% on average.

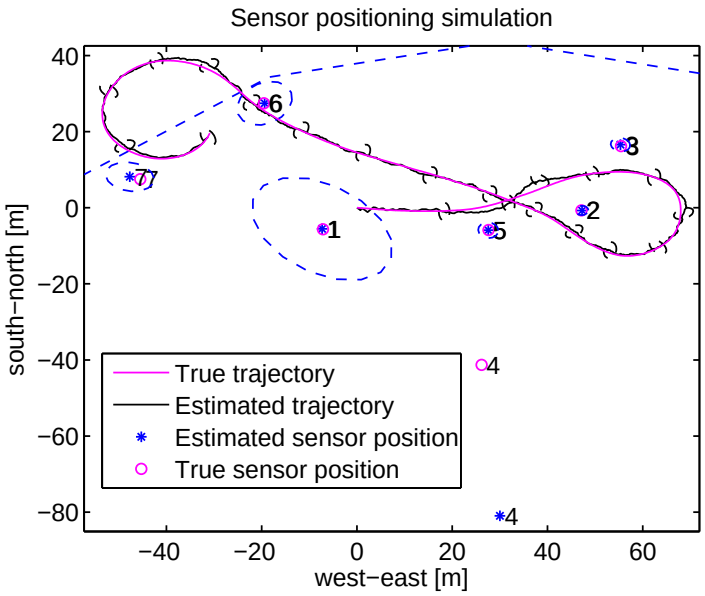


***Figure 5:*** *Estimated vessel trajectory and sensor positions with $2\sigma$ uncertainty. GNSSand magnetometers are used as sensors.*

## 3.3 Trajectory Evaluation using CRLB

CRLB for sensor positions surrounding a couple of trajectories were calculated for the case of GNSS and magnetometers. A high CRLB indicates that after a simulation, a sensor in that position would still have a high uncertainty. Fig. 7 shows the trajectory used in Sections 3.1 and 3.2. Fig. 8 and Fig. 9 show two other trajectories. It is clear that the CRLB becomes low in an area where the vessel can be observed from many directions. In Fig. 8 sensor positions quite close to the end of the trajectory have a high CRLB since they only observe the vessel from one direction. In Fig. 9 sensor positions between the start and end point of the trajectory are relatively difficult to estimate since it only observe the vessel from two opposite directions. The simulations suggest that in field experiments the vessel should be maneuvered in a trajectory that allows each sensor to observe the vessel from as many directions as possible.

## 3.4 Sensitivity Analysis, Magnetic Dipole

The magnetic dipole of the vessel will probably not be accurately measured in a real world experiment. How will the positioning perform if the estimated magnitude of the dipole is for example 102% or 110% of the true magnitude?

The trajectory previously used has been simulated using an assumed dipole that

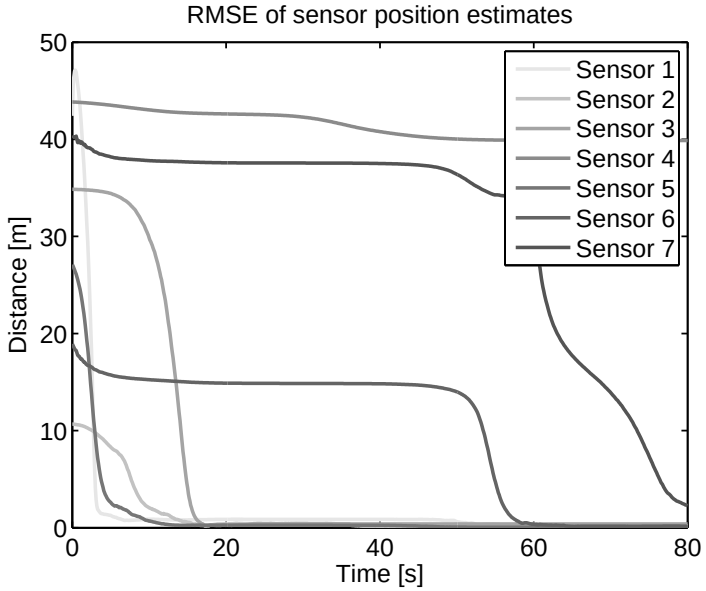


***Figure 6:*** *RMSE of estimated sensor positions throughout the simulation. GNSSand magnetometers are used as sensors.*

differs from the true one. A dipole with a magnitude of 98% of the true one is generated and the error is divided over the three components of the dipole. Each dipole error is simulated multiple times using the same trajectory and each time the error is distributed amongst the dipole components differently. Again, a sensor failing to retain the true sensor position within two standard deviations is considered incorrectly positioned. Table 2 shows the percentage of incorrectly positioned sensors for different errors of magnitude and different simulation settings.

| **Dipole** | 80% | 90% | 95% | 98% | 99% | 100% |
|---|---|---|---|---|---|---|
| SLAM | 44.6% | 25.7% | 23.4% | 23.4% | 18.9% | 14.3% |
| GNSS | 38.3% | 9.7% | 3.4% | 2.9% | 0.0% | 0.0% |
| **Dipole** | 101% | 102% | 105% | 110% | 120% | |
| SLAM | 14.3% | 14.3% | 16.6% | 34.3% | 53.1% | |
| GNSS | 4.0% | 4.6% | 8.6% | 12.0% | 38.3% | |

***Table 2:*** *Sensitivity analysis of error in dipole estimate.*

## 3.5 Sensitivity Analysis, Sensor Orientation

The sensor orientation is assumed measured in the previous experiments since it can be estimated prior to the experiment. We will now study how sensitive the system is to errors in the orientation estimate. The positioning performance

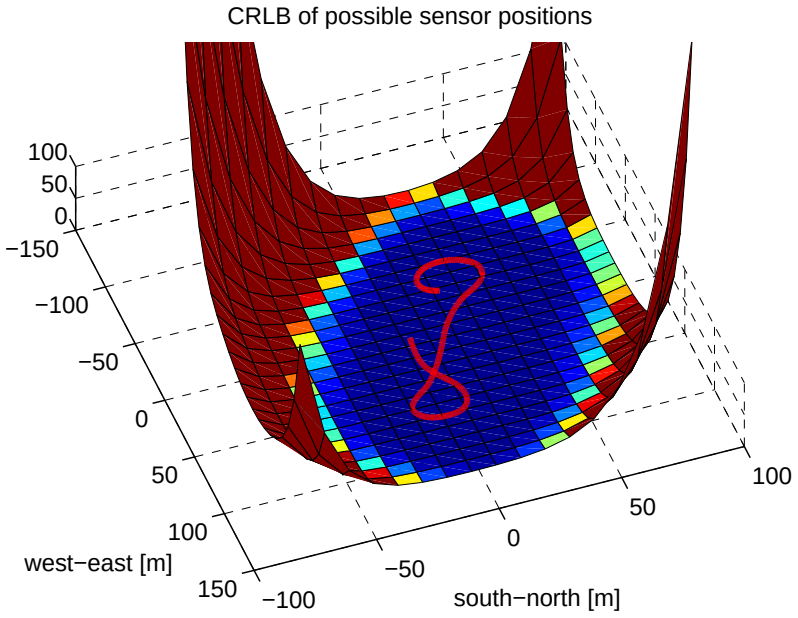


***Figure 7:*** *CRLB for all sensor positions surrounding the trajectory (in red). Trajectory 1.*

when sensor orientation errors are present, are evaluated using 25 Monte Carlo simulations for each orientation error using different trajectories. For each simulation, random orientation errors with the stated covariance are generated. (A covariance of 0.16 rad results in orientation errors up to 0.8 rad or 45°.) Table 3 shows the percentage of incorrectly positioned sensors for different sensor orientation error covariances.

Note that the sensor positioning error of a system using GNSS and magnetometers was merely unaffected by the introduction of an orientation covariance of up to 0.04 rad. If the sensor observes the vessel from many different directions, the positioning still succeeds. When only magnetometers are used, the trajectory measurements cannot compensate for the errors in orientation, rendering larger positioning errors.

| **Ori Cov** | 0.0 rad | 0.01 rad | 0.04 rad | 0.16 rad | 0.36 rad |
|---|---|---|---|---|---|
| SLAM | 26.3% | 29.8% | 36.9% | 54.8% | 52.4% |
| GNSS | 12.9% | 12.5% | 11.9% | 18.5% | 26.8% |

***Table 3:*** *Sensitivity analysis of error in estimated sensor orientation.*

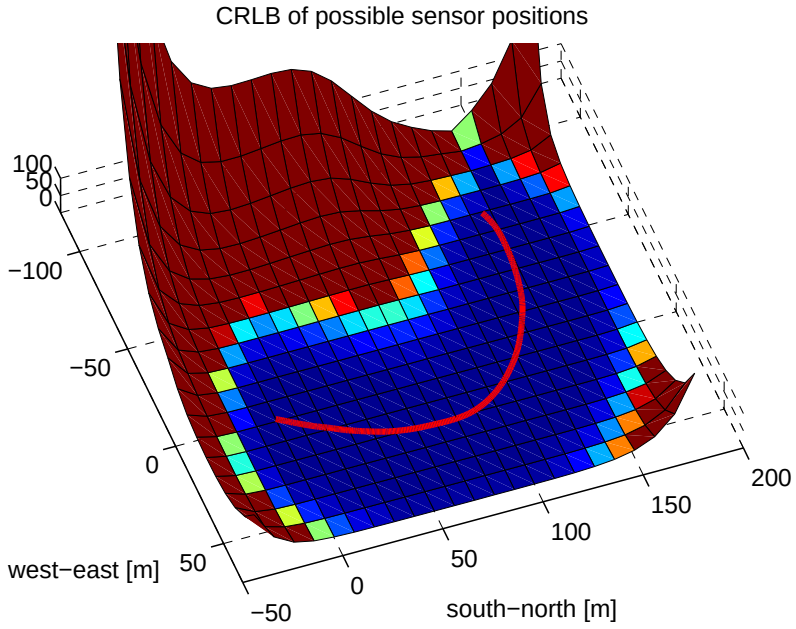


***Figure 8:*** *CRLB for all sensor positions surrounding the trajectory (in red). Trajectory 2.*

## 4 Conclusions

We have presented a silent underwater sensor localization scheme using triaxial magnetometers and a friendly vessel with known magnetic characteristics. More accurate sensor positions will enhance the detection, tracking and classification ability of the underwater sensor network. Monte Carlo simulations indicate that a sensor positioning accuracy of 26.3% is achievable when only magnetometers are used and of 12.9% when GNSS and magnetometers are used. Knowing the magnetic dipole of the vessel is important and a dipole magnitude error of 10% results in a positioning error increase of about 10%. Simulations also show that our positioning scheme is quite insensitive to minor errors in sensor orientation, when GNSS is used throughout the trajectory.

## Acknowledgments

This work was supported by the Strategic Research Center MOVIII, funded by the Swedish Foundation for Strategic Research, SSF, CADICS, a Linnaeus center funded by the Swedish Research Council, and LINK-SIC, an Industry Excellence Center founded by Vinnova.

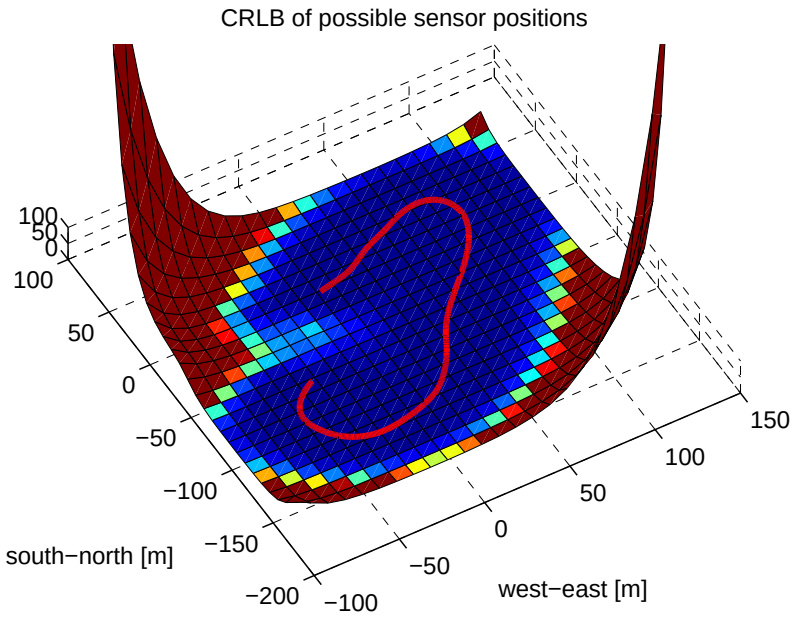


***Figure 9:*** *CRLB for all sensor positions surrounding the trajectory (in red). Trajectory 3.*

## Bibliography

I. F Akyildiz, D. Pompili, and T. Melodia. Underwater acoustic sensor networks: research challenges. *Elsevier Journal of Ad Hoc Networks*, 3(3), 2005.

J. Andrade-Cetto and A. Sanfeliu. The Effects of Partial Observability When Building Fully Correlated Maps. *Robotics, IEEE Transactions on*, 21(4):771–777, Aug. 2005. ISSN 1552-3098. doi: 10.1109/TRO.2004.842342.

T. Bailey and H. Durrant-Whyte. Simultaneous localization and mapping (SLAM): part II. *Robotics & Automation Magazine, IEEE*, 13(3):108 – 117, 2006.

M. Birsan. Non-linear kalman filters for tracking a magnetic dipole. In *Proc. of Intl. Conf. on Maritime Electromagnetics, MARELEC*, 2004.

M. Birsan. Unscented particle filter for tracking a magnetic dipole target. In *Proc. of MTS/IEE OCEANS*, 2005.

M. Birsan. Electromagnetic source localization in shallow waters using Bayesian matched-field inversion. *Inverse Problems*, 22(1):43–53, 2006.

M. Bryson and S. Sukkarieh. Observability analysis and active control for airborne SLAM. *Aerospace and Electronic Systems, IEEE Transactions on*, 44(1): 261–280, January 2008. ISSN 0018-9251. doi: 10.1109/TAES.2008.4517003.

J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010.

V. Chandrasekhar and W.KG. Seah. Area Localization Scheme for Underwater Sensor Networks. In *Proc. of the IEEE OCEANS Asia Pacific Conf.*, 2006.

V. Chandrasekhar, W.KG. Seah, Y. Sang Choo, and H.V. Ee. Localization in underwater sensor networks: survey and challenges. In *ACM Intl. Workshop on UnderWater Networks WUWNet*, 2006.

D.K. Cheng. *Field and Wave Electromagnetics*. Addison-Wesley, Reading, Massachusetts, 2nd edition, 1989.

W. Cheng, A.Y. Teymorian, L. Ma, X. Cheng, X. Lu, and Z. Lu. Underwater Localization in Sparse 3D Acoustic Sensor Networks. In *Proc. of conf. on Computer Communications, INFOCOM,IEEE*, 2008.

X. Cheng, H. Shu, Q. Liang, and DHC Du. Silent Positioning in Underwater Acoustic Sensor Networks. *Transactions on Vehicular Technology, IEEE*, 57(3), May. 2008.

E. Dalberg, A. Lauberts, R.K. Lennartsson, M.J. Levonen, and L. Persson. Underwater target tracking by means of acoustic and electromagnetic data fusion. In *Proc. of Intl. Conf. on Information Fusion*, 2006.

Z.A. Daya, D.L. Hutt, and T.C. Richards. Maritime electromagnetism and DRDC Signature Management research. Technical report, Defence R&D Canada, 2005.

F. Dellaert. Square root sam. In *Robotics: Science and Systems*, pages 177–184, 2005.

H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping (SLAM): part I. *Robotics & Automation Magazine, IEEE*, 13(2):99 – 110, 2006.

M. Erol, L.F.M. Vieira, and M. Gerla. AUV-Aided Localization for Underwater Sensor Networks. In *Proc. of Intl. Conf. on Wireless Algorithms, Systems and Applications, WASA*, 2007.

M. Erol, L.F.M. Vieira, A. Caruso, F. Paparella, M. Gerla, and S. Oktug. Multi Stage Underwater Sensor Localization using Mobile Beacons. In *Proc. of Intl. Conf. on Sensor Technologies and Applications, SENSORCOMM*, 2008.

T.I. Fossen and T. Perez. Marine Systems Simulator (MSS), 2004. www.marinecontrol.org.

F. Gustafsson. *Adaptive Filtering and Change Detection*. John Wiley & Sons, Hoboken, New Jersey, 2nd edition, 2001.

M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Fast Incremental Smoothing and Mapping with Efficient Data Association. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 1670–1677, April 2007. doi: 10.1109/ROBOT.2007.363563.

S. M. Kay. *Fundamentals of Statistical Signal Processing - Estimation Theory*. Prentice Hall, 1993.

J. Kim and S. Sukkarieh. Improving the real-time efficiency of inertial SLAM and understanding its observability. *Intelligent Robots and Systems, 2004. (IROS 2004). 2004 IEEE/RSJ International Conference on*, 1:21–26 vol.1, Oct. 2004. doi: 10.1109/IROS.2004.1389323.

K.W Lee, W.S. Wijesoma, and I.G. Javier. On the Observability and Observability Analysis of SLAM. In *Proceedings of the IEEE/RSJ International Conference Intelligent Robots and Systems*, 2006.

A. Lindin. Analysis and modelling of magnetic mine sweep for naval purposes. Master's thesis, Linköping University, 2007. [In swedish].

A. Lundin. Underwater electric signatures. Are they important for a future navy? Master's thesis, Swedish National Defence College, 2003. [In swedish].

J.B. Nelson and T.C Richards. Magnetic source parameters of MR OFFSHORE measured during trial MONGOOSE 07. Technical report, Defence R&D - Atlantic, Dartmouth NS (CAN), 2007.

L.D.L. Perera, A. Melkumyan, and E. Nettleton. On the linear and nonlinear observability analysis of the SLAM problem. *Mechatronics, 2009. ICM 2009. IEEE*

*International Conference on*, pages 1–6, April 2009. doi: 10.1109/ICMECH.2009.4957168.

C. Tian, W. Liu, J. Jin, Y. Wang, and Y. Mo. Localization and synchronization for 3D underwater acoustic sensor network. In *Intl. Conf. on Ubiquitous Intelligence and Computing UIC*, 2007.

Z. Wang and G. Dissanayake. Observability analysis of SLAM using fisher information matrix. *Control, Automation, Robotics and Vision, 2008. ICARCV 2008. 10th International Conference on*, pages 1242–1247, Dec. 2008. doi: 10.1109/ICARCV.2008.4795699.

S.Y. Wong, J.G. Lim, S.V. Rao, and W.KG. Seah. Multihop Localization with Density and Path Length Awareness in Non-Uniform Wireless Sensor Networks. In *Proc. of IEEE Vehicular Technology Conf.*, 2005.

Z. Zhou, J.-H. Cui, and S. Zhou. Localization for large-scale underwater sensor network. In *Proc. of IFIP Networking*, 2007.

**PhD Dissertations**
**Division of Automatic Control**
**Linköping University**

**M. Millnert:** Identification and control of systems subject to abrupt changes. Thesis No. 82, 1982. ISBN 91-7372-542-0.
**A. J. M. van Overbeek:** On-line structure selection for the identification of multivariable systems. Thesis No. 86, 1982. ISBN 91-7372-586-2.
**B. Bengtsson:** On some control problems for queues. Thesis No. 87, 1982. ISBN 91-7372-593-5.
**S. Ljung:** Fast algorithms for integral equations and least squares identification problems. Thesis No. 93, 1983. ISBN 91-7372-641-9.
**H. Jonson:** A Newton method for solving non-linear optimal control problems with general constraints. Thesis No. 104, 1983. ISBN 91-7372-718-0.
**E. Trulsson:** Adaptive control based on explicit criterion minimization. Thesis No. 106, 1983. ISBN 91-7372-728-8.
**K. Nordström:** Uncertainty, robustness and sensitivity reduction in the design of single input control systems. Thesis No. 162, 1987. ISBN 91-7870-170-8.
**B. Wahlberg:** On the identification and approximation of linear systems. Thesis No. 163, 1987. ISBN 91-7870-175-9.
**S. Gunnarsson:** Frequency domain aspects of modeling and control in adaptive systems. Thesis No. 194, 1988. ISBN 91-7870-380-8.
**A. Isaksson:** On system identification in one and two dimensions with signal processing applications. Thesis No. 196, 1988. ISBN 91-7870-383-2.
**M. Viberg:** Subspace fitting concepts in sensor array processing. Thesis No. 217, 1989. ISBN 91-7870-529-0.
**K. Forsman:** Constructive commutative algebra in nonlinear control theory. Thesis No. 261, 1991. ISBN 91-7870-827-3.
**F. Gustafsson:** Estimation of discrete parameters in linear systems. Thesis No. 271, 1992. ISBN 91-7870-876-1.
**P. Nagy:** Tools for knowledge-based signal processing with applications to system identification. Thesis No. 280, 1992. ISBN 91-7870-962-8.
**T. Svensson:** Mathematical tools and software for analysis and design of nonlinear control systems. Thesis No. 285, 1992. ISBN 91-7870-989-X.
**S. Andersson:** On dimension reduction in sensor array signal processing. Thesis No. 290, 1992. ISBN 91-7871-015-4.
**H. Hjalmarsson:** Aspects on incomplete modeling in system identification. Thesis No. 298, 1993. ISBN 91-7871-070-7.
**I. Klein:** Automatic synthesis of sequential control schemes. Thesis No. 305, 1993. ISBN 91-7871-090-1.
**J.-E. Strömberg:** A mode switching modelling philosophy. Thesis No. 353, 1994. ISBN 91-7871-430-3.
**K. Wang Chen:** Transformation and symbolic calculations in filtering and control. Thesis No. 361, 1994. ISBN 91-7871-467-2.
**T. McKelvey:** Identification of state-space models from time and frequency data. Thesis No. 380, 1995. ISBN 91-7871-531-8.
**J. Sjöberg:** Non-linear system identification with neural networks. Thesis No. 381, 1995. ISBN 91-7871-534-2.
**R. Germundsson:** Symbolic systems – theory, computation and applications. Thesis No. 389, 1995. ISBN 91-7871-578-4.

**P. Pucar:** Modeling and segmentation using multiple models. Thesis No. 405, 1995. ISBN 91-7871-627-6.
**H. Fortell:** Algebraic approaches to normal forms and zero dynamics. Thesis No. 407, 1995. ISBN 91-7871-629-2.
**A. Helmersson:** Methods for robust gain scheduling. Thesis No. 406, 1995. ISBN 91-7871-628-4.
**P. Lindskog:** Methods, algorithms and tools for system identification based on prior knowledge. Thesis No. 436, 1996. ISBN 91-7871-424-8.
**J. Gunnarsson:** Symbolic methods and tools for discrete event dynamic systems. Thesis No. 477, 1997. ISBN 91-7871-917-8.
**M. Jirstrand:** Constructive methods for inequality constraints in control. Thesis No. 527, 1998. ISBN 91-7219-187-2.
**U. Forssell:** Closed-loop identification: Methods, theory, and applications. Thesis No. 566, 1999. ISBN 91-7219-432-4.
**A. Stenman:** Model on demand: Algorithms, analysis and applications. Thesis No. 571, 1999. ISBN 91-7219-450-2.
**N. Bergman:** Recursive Bayesian estimation: Navigation and tracking applications. Thesis No. 579, 1999. ISBN 91-7219-473-1.
**K. Edström:** Switched bond graphs: Simulation and analysis. Thesis No. 586, 1999. ISBN 91-7219-493-6.
**M. Larsson:** Behavioral and structural model based approaches to discrete diagnosis. Thesis No. 608, 1999. ISBN 91-7219-615-5.
**F. Gunnarsson:** Power control in cellular radio systems: Analysis, design and estimation. Thesis No. 623, 2000. ISBN 91-7219-689-0.
**V. Einarsson:** Model checking methods for mode switching systems. Thesis No. 652, 2000. ISBN 91-7219-836-2.
**M. Norrlöf:** Iterative learning control: Analysis, design, and experiments. Thesis No. 653, 2000. ISBN 91-7219-837-0.
**F. Tjärnström:** Variance expressions and model reduction in system identification. Thesis No. 730, 2002. ISBN 91-7373-253-2.
**J. Löfberg:** Minimax approaches to robust model predictive control. Thesis No. 812, 2003. ISBN 91-7373-622-8.
**J. Roll:** Local and piecewise affine approaches to system identification. Thesis No. 802, 2003. ISBN 91-7373-608-2.
**J. Elbornsson:** Analysis, estimation and compensation of mismatch effects in A/D converters. Thesis No. 811, 2003. ISBN 91-7373-621-X.
**O. Härkegård:** Backstepping and control allocation with applications to flight control. Thesis No. 820, 2003. ISBN 91-7373-647-3.
**R. Wallin:** Optimization algorithms for system analysis and identification. Thesis No. 919, 2004. ISBN 91-85297-19-4.
**D. Lindgren:** Projection methods for classification and identification. Thesis No. 915, 2005. ISBN 91-85297-06-2.
**R. Karlsson:** Particle Filtering for Positioning and Tracking Applications. Thesis No. 924, 2005. ISBN 91-85297-34-8.
**J. Jansson:** Collision Avoidance Theory with Applications to Automotive Collision Mitigation. Thesis No. 950, 2005. ISBN 91-85299-45-6.
**E. Geijer Lundin:** Uplink Load in CDMA Cellular Radio Systems. Thesis No. 977, 2005. ISBN 91-85457-49-3.
**M. Enqvist:** Linear Models of Nonlinear Systems. Thesis No. 985, 2005. ISBN 91-85457-64-7.
**T. B. Schön:** Estimation of Nonlinear Dynamic Systems — Theory and Applications. Thesis No. 998, 2006. ISBN 91-85497-03-7.

**I. Lind:** Regressor and Structure Selection — Uses of ANOVA in System Identification. Thesis No. 1012, 2006. ISBN 91-85523-98-4.

**J. Gillberg:** Frequency Domain Identification of Continuous-Time Systems Reconstruction and Robustness. Thesis No. 1031, 2006. ISBN 91-85523-34-8.

**M. Gerdin:** Identification and Estimation for Models Described by Differential-Algebraic Equations. Thesis No. 1046, 2006. ISBN 91-85643-87-4.

**C. Grönwall:** Ground Object Recognition using Laser Radar Data – Geometric Fitting, Performance Analysis, and Applications. Thesis No. 1055, 2006. ISBN 91-85643-53-X.

**A. Eidehall:** Tracking and threat assessment for automotive collision avoidance. Thesis No. 1066, 2007. ISBN 91-85643-10-6.

**F. Eng:** Non-Uniform Sampling in Statistical Signal Processing. Thesis No. 1082, 2007. ISBN 978-91-85715-49-7.

**E. Wernholt:** Multivariable Frequency-Domain Identification of Industrial Robots. Thesis No. 1138, 2007. ISBN 978-91-85895-72-4.

**D. Axehill:** Integer Quadratic Programming for Control and Communication. Thesis No. 1158, 2008. ISBN 978-91-85523-03-0.

**G. Hendeby:** Performance and Implementation Aspects of Nonlinear Filtering. Thesis No. 1161, 2008. ISBN 978-91-7393-979-9.

**J. Sjöberg:** Optimal Control and Model Reduction of Nonlinear DAE Models. Thesis No. 1166, 2008. ISBN 978-91-7393-964-5.

**D. Törnqvist:** Estimation and Detection with Applications to Navigation. Thesis No. 1216, 2008. ISBN 978-91-7393-785-6.

**P-J. Nordlund:** Efficient Estimation and Detection Methods for Airborne Applications. Thesis No. 1231, 2008. ISBN 978-91-7393-720-7.

**H. Tidefelt:** Differential-algebraic equations and matrix-valued singular perturbation. Thesis No. 1292, 2009. ISBN 978-91-7393-479-4.

**H. Ohlsson:** Regularization for Sparseness and Smoothness — Applications in System Identification and Signal Processing. Thesis No. 1351, 2010. ISBN 978-91-7393-287-5.

**S. Moberg:** Modeling and Control of Flexible Manipulators. Thesis No. 1349, 2010. ISBN 978-91-7393-289-9.

**J. Wallén:** Estimation-based iterative learning control. Thesis No. 1358, 2011. ISBN 978-91-7393-255-4.

**J. Hol:** Sensor Fusion and Calibration of Inertial Sensors, Vision, Ultra-Wideband and GPS. Thesis No. 1368, 2011. ISBN 978-91-7393-197-7.

**D. Ankelhed:** On the Design of Low Order H-infinity Controllers. Thesis No. 1371, 2011. ISBN 978-91-7393-157-1.

**C. Lundquist:** Sensor Fusion for Automotive Applications. Thesis No. 1409, 2011. ISBN 978-91-7393-023-9.

**P. Skoglar:** Tracking and Planning for Surveillance Applications. Thesis No. 1432, 2012. ISBN 978-91-7519-941-2.

**K. Granström:** Extended target tracking using PHD filters. Thesis No. 1476, 2012. ISBN 978-91-7519-796-8.

**C. Lyzell:** Structural Reformulations in System Identification. Thesis No. 1475, 2012. ISBN 978-91-7519-800-2.